KNOWLEDGE-BASED JUDGEMENTS OF CAUSALITY

KNOWLEDGE-BASED JUDGEMENTS OF CAUSALITY: CONTIGUITY, CONGRUITY, AND DIRECTION OF THE CAUSAL ARROW

By
JASON M. TANGEN, BASc

A Dissertation
Submitted to the School of Graduate Studies
in Partial Fulfillment of the Requirements
for the Degree
Doctor of Philosophy

McMaster University

DOCTOR OF PHILOSOPHY (2003)
(Psychology)

McMaster University
Hamilton, Ontario

TITLE: Knowledge-based judgements of causality: Contiguity, congruity, and direction of the causal arrow

AUTHOR: Jason M. Tangen, BASc (University of Lethbridge)

SUPERVISOR: Professor L. G. Allan

NUMBER OF PAGES: xiii, 122

Psychology Press
Taylor & Francis Group plc

To: Jason Tangen
Department of Psychology
1280 Main Street W
Hamilton, ON, Canada
L8S 4K1

From: Isobel Muir, Permissions, Psychology Press

Date: 12th June 2003

Re: Jason M. Tangen and Lorraine G. Allan (2003) The relative effect of cue interaction, *Quarterly Journal of Experimental Psychology, 56(B)* pp. 279 – 300.

Permission granted, free of charge, for the use specified in your original request only. Proper credit must be given to our publication, including: title; author(s) and/or editor(s); copyright date; and the words " Reprinted by permission of The Experimental Psychology Society."

Isobel Muir
Journals Editorial Assistant
PSYCHOLOGY PRESS LIMITED
Isobel.muir@psypress.co.uk

Psychology Press Ltd
27 Church Road
Hove, East Sussex BN3 2FA
Tel: +44 (0) 1273 207411
Fax: +44 (0) 1273 205612
www.psypress.co.uk

Distribution (Journals)
Rankine Road
Basingstoke
Hampshire RG24 8PR
Tel: +44 (0) 1256 813000
Fax: +44 (0) 1256 479438

Distribution (Books)
Cheriton House
North Way, Andover
Hampshire SP10 5BE
Tel: +44 (0) 1264 332424
Fax: +44 (0) 1264 364418

Registered Office 11 New Fetter Lane London EC4P 4EE. Registered No.1493044

ansaction Books • Periodicals • Fulfillment & Distribution • International Shipping

INVOICE NO.: 97927

Jason Tangen
Department of Psychology
1280 Main Street West
Hamilton, Ontario
Canada L8S 4K1

Rutgers University
35 Berrue Circle
Piscataway, NJ 08854-8042
Telephone 732 / 445-2280
Facsimile 732 / 445-3138
E-mail trans@transactionpub.com
Website http://www.transactionpub.com

Date: June 2003
Reference:
Payable prior to first distribution

Permission is granted to reproduce the following material subject to publication of acknowledgment to Transaction Publishers and the copyright holder, and payment of the fee shown:

Author:	**J.M. Tangen, Allan, L.G., R. Wood, and T. Shah**
Article/Chapter Title:	**Temporal Contiguity and Contingency Judgments**
Book/Journal Title:	**Integrative Physiological and Behavior Science**
Volume No.:	
Issue No.:	
Copyright Date:	**19**
Page Reference:	
Maximum Number of Copies:	
Amount due:	**$N/C** (PERMISSION IS GRANTED FOR REPRINT PURPOSES ONLY)

ED ID #430908220

Terms of License:

MARKET: Permission granted is non-exclusive and, unless otherwise stated, is valid throughout the world in hard copy form only, and in English language only.

ACKNOWLEDGEMENT: A credit line in the form indicated must appear on the title page of the article/chapter/excerpt of one page or more, and directly beneath the exhibit or short excerpt with applicable title, author, publication date and copyright.

This permission does not allow the use of our material in any other edition, or by any other means of reproduction. Separate permission request must be received by Transaction for any such additional use.

Reprinted by permission of Transaction Publishers. "Article/chapter" title by (author, issue/month, year). Copyright @ (yr.) by Transaction Publishers.

Marlèna Davidian
Permissions Manager

Abstract

Under what circumstances does knowledge of causal asymmetry and temporal delay influence causal judgements? We begin a series of thirteen experiments by providing evidence that both high-level (causal reasoning) processes, and low-level (associative) processes influence causal assessment depending on what is asked about the events. Specifically, participants were more sensitive to causal structure in their ratings than in their prediction responses, on earlier rather than later trials, and when asked to provide an integrative causal rating. Emphasising the direction and nature of the causal relationship and the wording of the test question had no influence on participants' sensitivity to causal asymmetry. Next, we provide evidence that participants' ratings track conditional rather than uncondtional contingencies as predicted by the conditional ΔP account as well as the Rescorla-Wagner model at asymptote. Our results suggest that participants tend to rate the influence of each cause conditional on the absence of the other cause. This tendency is not reflected by the Rescorla-Wagner model. Finally, we examine the role of temporal contiguity on judgments of contingency using a human analogue of the Pavlovian task. Our results suggest that knowledge of temporal delay modulates causal judgements.

Acknowledgements

I am sincerely grateful to my supervisor, Lorraine Allan, for her thoughtful advice, patience, and friendship. She has provided me with incredible support, inspiration, and guidance in both the writing of this dissertation and the research that went into it. Her confidence in my ability to do meaningful research and work independently, has been a most important source of encouragement. I would like to thank Lee Brooks who has given me gifts of kindness and wisdom. I appreciate our numerous discussions, which helped guide and shape my research and future work. Lee is a true scholar, a generous and conscientious mentor. I also want to acknowledge John Vokey who is solely responsible for my pursuing a career in academia, and for the skills, knowledge, and motivation to excel in graduate school. Not only is he a fantastic undergraduate supervisor, but he is also one of my closest friends. Thanks for everything John. Without question, these three have taken it upon themselves to ensure my happiness and success. I can never repay what you have given me. I can only do my best to pay it forward to my own students.

I also want to thank Sue Becker for her careful consideration and constructive suggestions. I would also like to express my appreciation to Ralph Miller as my external examiner, for his thoughtful questions at my defence.

Financial support was provided by NSERC and OGS. Thanks to Victor Galleguillos, Taral Shah, Hedyeh Sadeghi, Erin Wilson, Robert Wood, Meghan Burke, Jim Provost, and Christine Moore for their help in running experiments and analysing data. Thanks also to the Psychology administrative staff: Milica Pavlica, Wendy Selbie, Sally Thompson, Nancy Riddell, and above all Erie Long. Thanks also to the creators of LaTeX and TeXShop, making my dissertation writing process untroublesome and enjoyable.

I am grateful to my friends and colleagues: Alex Ophir, Kevin Eva, Betsy Agar, Andrew Clark, Brett Beston, Nicole Anderson, Matt Crump, Sam Hannah, and Vito Scavetta for their support and encouragement during my time at McMaster.

I would like to thank my family for their unwavering support and encouragement, who always told me that I could accomplish anything that I set my mind to and who taught me discipline, respect for others, and to finish what I started.

But despite all of these significant contributions, this dissertation might never have come to be if it were not for Melanie McKenzie. She has tolerated my impatience, soothed my irritation, and been there when I needed her the most. She has spent hours listening to my talks and providing feedback and criticism. She has put up with me, pushed me, and made me laugh. Thank you Melanie.

Table of Contents

Abstract iii

Acknowledgements iv

List of Figures viii

List of Tables x

Preface xii

1 Introduction 1
- 1.1 Blocking Effects and Associations . . . 2
- 1.2 Blocking Effects in Human Learning . . . 4
  - 1.2.1 Causal-Model Theory . . . 5

2 Cue-interaction and Judgments of Causality: Contributions of Causal and Associative Processes 6
- 2.1 Preface . . . 6
- 2.2 Abstract . . . 6
- 2.3 Introduction . . . 7
- 2.4 Experiment 1 . . . 10
  - 2.4.1 Method . . . 11
- 2.5 Experiment 2 . . . 16
  - 2.5.1 Method . . . 17
- 2.6 Experiment 3 . . . 20
  - 2.6.1 Method . . . 21
- 2.7 Experiment 4 . . . 26
  - 2.7.1 Method . . . 26
- 2.8 General Discussion . . . 34

3 Assessing (in)sensitivity to causal asymmetry: A matter of degree 37
- 3.1 Preface . . . 37
- 3.2 Abstract . . . 37

3.3 Introduction . . . 38
3.4 Causal-Model Theory and Conditional ΔP . . . 40
3.5 Alleged Causal Model Influences . . . 42
3.5.1 General Method . . . 42
3.6 Experiment 1: Clarity of the Causal Model . . . 45
3.6.1 Method . . . 45
3.6.2 Results . . . 46
3.6.3 Discussion . . . 49
3.7 Experiment 2: Test Question . . . 49
3.7.1 Method . . . 49
3.7.2 Results . . . 50
3.7.3 Discussion . . . 54
3.8 Experiment 3: Integration . . . 54
3.8.1 Method . . . 55
3.8.2 Results . . . 56
3.8.3 Discussion . . . 59
3.9 General Discussion . . . 59

4 The Relative Effect of Cue-interaction **61**
4.1 Preface . . . 61
4.2 Abstract . . . 62
4.3 Introduction . . . 62
4.4 Method . . . 69
4.4.1 Observers . . . 69
4.4.2 Apparatus . . . 69
4.4.3 Procedure . . . 69
4.5 Results . . . 70
4.5.1 Rating Data . . . 70
4.5.2 Prediction Data . . . 72
4.5.3 The Rescorla-Wagner Model . . . 76
4.6 Discussion . . . 80
4.7 Postscript . . . 81
4.8 Method . . . 82
4.8.1 Design . . . 82
4.9 Results . . . 82

5 Temporal Contiguity and Contingency Judgments: A Pavlovian Analogue **85**
5.1 Preface . . . 85
5.2 Abstract . . . 86
5.3 Introduction . . . 86
5.4 Experiment 1 . . . 88
5.4.1 Method . . . 88
5.4.2 Results and Discussion . . . 89

5.5 Experiment 2 91
5.5.1 Method 92
5.5.2 Results and Discussion 93
5.5.3 Discussion 95

6 General Discussion **97**
6.0.4 Shortcomings and Future Directions 102
6.1 Conclusion 104

References **105**

A Causal Model Prompts **112**
A.0.1 2C-1E 112
A.0.2 2E-1C 113

B Instructions for Chapter 4 **114**

C Rescorla-Wagner = Conditional ΔP **116**

D Instructions for Chapter 5 **120**
D.1 Instructions for Experiment 1 120
D.2 Instructions for Experiment 2 121
D.2.1 Immediate Instructions - Remote-Control Detonator 121
D.2.2 Delay Instructions - Grenade Launcher 122

List of Figures

2.1 Summary 4×2 contingency matrix illustrating each of the possible cause-effect combinations for two cues. Each cell represents the frequency of each event type. 9

2.2 Four possible causal scenarios generated by crossing causal order (CE vs. EC) with the number of cues and outcomes (2-1 vs. 1-2). 10

2.3 Mean ratings in Experiment 1 after 48 trials of Event A (Figure 2.3a) and of Event B (Figure 2.3b). For each event, the ratings are shown as a function of ΔP_B separately for each of the four scenarios. Error-bars represent standard errors of the means. 14

2.4 Mean ratings in Experiment 2 after 32 trials of Event A (Figure 2.4a) and of Event B (Figure 2.4b). For each event, the ratings are shown as a function of ΔP_B (0, 0.25, 0.5, 0.75, 1) separately for each of the four scenarios. Error-bars represent standard errors of the means. 18

2.5 Mean ratings in Experiment 3 after 32 trials of Cue A (Figure 2.5a) and of Cue B (Figure 2.5b). For each cue, the ratings are shown as a function of ΔP_B (0, 0.25, 0.5, 0.75, 1) separately for each of the two scenarios. Error-bars represent standard errors of the means. 23

2.6 Mean estimated conditional ΔP values in Experiment 3 for Cues A and B as a function of ΔP_B separately for the two causal scenarios. $est\Delta P_{A|B}$ is shown in Figure 2.6a, $est\Delta P_{A|\sim B}$ is shown in Figure 2.6b, $est\Delta P_{B|A}$ is shown in Figure 2.6c, and $est\Delta P_{B|\sim A}$ is shown in Figure 2.6d. Error-bars represent standard errors of the means. 24

2.7 Mean ratings in Experiment 4 after 64 trials of Event A (Figure 2.7a) and of Event B (Figure 2.7b). For each cue, the ratings are shown as a function of ΔP_B (0, 0.25, 0.75, 1) separately for each of the two conditions. Error-bars represent standard errors of the means. 28

2.8 Mean estimated conditional ΔP values in Experiment 4 for Events A and B as a function of ΔP_B separately for the two causal scenarios after 64 trials. $est\Delta P_{A|B}$ is shown in Figure 2.8a, $est\Delta P_{A|\sim B}$ is shown in Figure 2.8b, $est\Delta P_{B|A}$ is shown in Figure 2.8c, and $est\Delta P_{B|\sim A}$ is shown in Figure 2.8d. Error-bars represent standard errors of the means. 32

3.1 Generic casual structure among three interconnected events. A1, A2, and A3 each represent a cause or an effect. 39

3.2 4×2 contingency matrix illustrating the eight possible cue-outcome combinations for two cues. Each cell represents the frequency of each event type. 40
3.3 Mean ratings in Experiment 1 after 64 trials of Cue A (Figure 3.3). The ratings are shown as a function of ΔP_B (0, 0.25, 0.75, 1) separately for each of the two conditions. Error-bars represent standard errors of the means. 46
3.4 Mean ratings in Experiment 2 after 64 trials of Cue A for participants given the *Causal* test question (Figure 3.4a) and the *Contiguous* test question (Figure 3.4b). The ratings are shown as a function of ΔP_B (0, 0.25, 0.75, 1) separately for each of the two conditions. Error-bars represent standard errors of the means. 51
3.5 Mean ratings in Experiment 3 after 64 trials of Cue A (Figure 3.5). The ratings are shown as a function of ΔP_B (0, 0.25, 0.75, 1) separately for each of the two conditions. Error-bars represent standard errors of the means. 56
4.1 Summary 2×2 contingency matrix illustrating each of the possible cause-effect combinations for one cue and one outcome. Each cell represents the frequency of each event type. 63
4.2 Summary 4×2 contingency matrix for the single-phase simultaneous blocking task. The notation in Figure 4.2a is for the conditional ΔP account, and the notation in Figure 4.2b is for the Rescorla-Wagner model. 64
4.3 Mean ratings for each condition (Baseline, Equal, and Unequal) are plotted as a function of trial (36, 54, and 72). Figure 4.3a presents the data for Cue A and Figure 4.3b presents the data for Cue B. 71
4.4 Mean estimated conditional ΔP values for Cue A are plotted as a function of trial (36, 54, and 72) for each condition (Baseline, Equal, and Unequal). $est\Delta P_{A|B}$ is shown in Figure 4.4a and $est\Delta P_{A|\sim B}$ is shown in Figure 4.4b. . . . 73
4.5 Mean estimated conditional ΔP values for Cue B are plotted as a function of trial (36, 54, and 72) for each condition (Baseline, Equal, and Unequal). $est\Delta P_{B|A}$ is shown in Figure 4.5a and $est\Delta P_{B|\sim A}$ is shown in Figure 4.5b. . . . 75
4.6 Rescorla-Wagner simulations for the three conditions (Baseline, Equal, and Unequal) are plotted in blocks of 10 trials. The predictive strength of Cue A (V_A) is shown in Figure 4.6a, of Cue B (V_B) in Figure 4.6b, and of the context (V_X) in Figure 4.6c. 77
4.7 Rescorla-Wagner simulations for the three conditions (Baseline, Equal, and Unequal) are plotted for the first 72 trials. The predictive strength of Cue A (V_A) is shown in Figure 4.7a, of Cue B (V_B) in Figure 4.7b, and of the context (V_X) in Figure 4.7c. 79
5.1 Mean ratings in Experiment 1 as a function of ΔP for each delay (unfilled bars for .4 sec and filled bars for 2.0 sec). The ratings at trial 20 are seen in Figure 5.1a and the ratings for trial 40 are seen in Figure 5.1b. 90
5.2 Mean ratings in Experiment 2 as a function of delay for the two cover stories (triangles for IC and squares for DC) at each ΔP value (filled symbols for $\Delta P = 0$ and unfilled symbols for $\Delta P = .5$). The ratings for the 0/5 order are in Figure 5.2a and the ratings for the 5/0 order are in Figure 5.2b. 94

List of Tables

1.1 Experimental Design by Kamin (1968) 2
1.2 Experimental Design by Wagner et al. (1968) 3
2.1 Frequency of events in Experiment 1. The unconditional ΔP values were calculated using Equations 2.1 and 2.2. The conditional ΔP values were calculated using Equations 2.3-2.6. 11
2.2 Frequency of events in Experiments 2 and 3. Unconditional ΔP values were calculated using Equations 2.1 and 2.2. Conditional ΔP values were calculated using Equations 2.3-2.6. 16
2.3 Frequency of events in Experiment 4. Unconditional ΔP values were calculated using Equations 2.1 and 2.2. Conditional ΔP values were calculated using Equations 2.3-2.6. 27
2.4 Experiment 4 overall ratings and estimated conditional ΔP values for Cue A after 32, 48, and 64 trials. 29
2.5 Experiment 4 overall ratings and estimated conditional ΔP values for Cue B after 32, 48, and 64 trials. 31
3.1 Frequency of events in Experiment 1-3. Unconditional ΔP values were calculated using Equations 3.1 and 3.2. Conditional ΔP values were calculated using Equations 3.3-3.6. 43
3.2 Experiment 1 mean ratings and estimated ΔP values of Cue A conditional on the presence of Cue B ($est\Delta P_{A|B}$) and the absence of Cue B ($est\Delta P_{A|\sim B}$) after 32, 48, and 64 trials. 47
3.3 Experiment 2 mean ratings and estimated ΔP values of Cue A conditional on the presence of Cue B ($est\Delta P_{A|B}$) and the absence of Cue B ($est\Delta P_{A|\sim B}$) after 32, 48, and 64 trials for the Causal test question. 52
3.4 Experiment 2 mean ratings and estimated ΔP values of Cue A conditional on the presence of Cue B ($est\Delta P_{A|B}$) and the absence of Cue B ($est\Delta P_{A|\sim B}$) after 32, 48, and 64 trials for the Contiguous test question. 53
3.5 Experiment 3 mean ratings and estimated ΔP values of Cue A conditional on the presence of Cue B ($est\Delta P_{A|B}$) and the absence of Cue B ($est\Delta P_{A|\sim B}$) after 32, 48, and 64 trials. 57
3.6 Experiment 3 integrative ratings of Cue A. 58

4.1 Frequency matrices, and unconditional and conditional probabilities for the situation where a moderately predictive cue is paired with a perfect predictor (.5/1) and for the situation where a moderately predictive cue is paired with a non predictive cue (.5/0). 65
4.2 4×2 matricies showing the frequencies suggested by Spellman (1996b, Figure 8) for three pairings of two cues. 67
4.3 Frequency of events in Experiment 1-3, mean ratings, actual and estimated (un)conditional ΔP values after 72 trials. The Rescorla-Wagner simulations were based on the same parameter values as indicated in Chapter 4, 360 trials, and 100 iterations. 83
5.1 Standard 2 × 2 contingency matrix for the human analogue of the Pavlovian task 87
5.2 Cell frequencies and conditional probabilities for Experiments 1 and 2 88

Preface

This dissertation is divided into four main sections. Each section represents either published material (Tangen & Allan, 2003; Allan, Tangen, Wood, & Shah, in press), material submitted for publication (Tangen & Allan, submitted) or to be submitted (Tangen, Allan, & Sadeghi, invited). Because these articles have multiple authorship, my contribution to each is explained here.

Tangen and Allan (submitted) report four experiments. All four experiments represent my contributions to the paper, and thus are all relevant to this dissertation. We provide evidence that both high and low-level processing influence causal assessment depending on what is being asked about the events, and participants' experience with those events. These data provide a novel contribution to the literature. Experiments 1-4 were conducted between January 2001 and October 2002.

Tangen et al. (invited) is an invited chapter stemming from the 2003 Associative Learning Symposium in Gregynog, Wales, UK. I was the first author on the presented paper, and therefore the first author on the published paper. All three experiments represent my contribution to the paper. Experiments 1 and 3 were conducted as part of an undergraduate research project by the latter author. The data provide an extension of the work presented in Tangen and Allan (submitted) by investigating the predictions made by our dual-process model. The work provides a novel contribution. The data from Experiment 1 was collected in January 2002, and Experiments 2 and 3 were conducted in November 2002 and January 2003 respectively.

Tangen and Allan (2003) report one experiment and simulation data pertaining to the three conditions examined in the paper. Both the experiment and simulation represent my contribution to the paper, and thus are relevant for this dissertation. We compare the Rescorla-Wagner model to conditional ΔP and derive an algebraic derivation equating the two models. As such, this paper constitutes an original contribution to the literature. The data from Experiment 1 were collected in January 2002.

The three experiments presented in the postscript of Chapter 4 were conducted as part of an undergraduate research project by Erin Wilson. The results from these experiments suggest that participants tend to rate the influence of each cause conditional on the absence of the other cause. The data were collected in October 2002.

Allan et al. (in press) is an invited paper resulting from the 2002 meeting of the Pavlovian Society in Westwood, California. Dr. Lorraine Allan was the first author on the presented paper, and therefore the first author on the published paper. Both experiments represent

my contribution to the paper. This research was conducted as part of two undergraduate research projects by the two latter authors. The two experiments examine the role of temporal contiguity on judgments of contingency. Unlike previous investigations that used instrumental conditioning procedures, we used a human analogue of the Pavlovian task. The data show that the effect of the actual delay on contingency judgment depends on the observer's expectation regarding the delay. This work provides an original contribution to the literature. The data from Experiment 1 and 2 of this paper were collected in November 2001 and October 2002 respectively.

Chapter 1

Introduction

> Our minds can shape the way a thing will be because we act according to our expectations.
>
> *Federico Fellini*

When confronted with a number of novel event relationships, under what circumstances does our general knowledge of causal direction and temporal delay guide judgements of causality? For example, if we asked Ruth to assess how predictive two viruses are in causing a disease, would her assessment of each virus differ from that of the two symptoms that Arthur is assessing as diagnostic of a particular illness? Does the fact that "Ruth is predicting" and "Arthur is diagnosing" influence how they regard each virus/symptom? If Ruth expects a long delay between the exposure to the virus and the onset of a symptom, when in fact, they occur simultaneously, how will her assessment compare to Arthur's who expects a short delay and observes a long delay? Will their assessments differ as the causal strength of each event changes? In this dissertation, I will focus primarily on the direction of the causal arrow. How knowledge of causal asymmetry affects judgements of causality.

I begin by describing the primary research and methodology behind the question of causal asymmetry. Chapter 2 describes four experiments that investigate the extent to which the causal structure of events influences causal judgements. As described below, sensitivity to causal asymmetry is measured by cue-interaction: the extent to which participants judge one event in light of another. The data from these experiments provide evidence that participants are sensitive to both the causal and temporal structure of the events depending on what and when they are asked about the events. Chapter 3 presents three experiments that test various hypotheses which makes predictions about the circumstances under which causal asymmetry affects judgements of causality. Chapter 4 provides a detailed investigation of the mechanism of cue-interaction. Specifically, in assessing the influence of two cues and one outcome, how do judgements differ between cues that vary according to their unconditional or conditional contingencies. Three related experiments are presented in the postscript of Chapter 4. Finally, Chapter 5 examines the role of temporal knowledge on judgements of causality. In particular, the congruity between observed and expected temporal delays is investigated. The results

Table 1.1: Experimental Design by Kamin (1968)

	Phase 1	Phase 2	Test
Group 1	A → O	AB → O	B ↓
Group 2		AB → O	B ↑

from each of the thirteen experiments in this dissertation are discussed in relation to one another under the broad scope of knowledge-based judgements of causality.

1.1 Blocking Effects and Associations

Contiguity or the pairing of events has long been recognised by learning theorists as insufficient to explain basic associative processes. In 1968, Leon Kamin described the blocking phenomenon as a demonstration of this insufficiency (Kamin, 1968). Using a two-phase design, as shown in Table 1.1, Kamin conditioned a group of animals to associate a single cue with an outcome (i.e., A → O). In Phase 2, a second cue was paired alongside the first (i.e., AB → O). A second group of animals were only exposed to the latter phase. Despite the extensive pairing of Cue B and the outcome, Group 1 learned very little about B compared to Group 2 (and other relevant control groups). The initial training with A blocks conditioning to the superimposed cue resulted in an attenuated response to B at test.

At the same time, Wagner, Logan, Haberlandt, and Price (1968) demonstrated a related phenomenon they labelled *relative validity* to demonstrate the same insufficiency of contiguity. Table 1.2 illustrates the design used by Wagner and his colleagues. They exposed animals to two compounds containing a common A cue that was paired either with Cue B or C. The animals were assigned to one of two groups: In Group 1, the AB compound was always paired with the outcome (100%) while the AC compound was never paired with the outcome (0%). In Group 2, each compound was paired with the outcome on 50% of the trials. During the test phase, the animals responded less to Cue A in Group 1 than in Group 2, even though it was paired with the outcome on 50% of the trials in both groups. These results demonstrate that the animals were not sensitive to the absolute validity of each cue, but rather were sensitive to the validity of each cue relative to one another.

The blocking and relative validity effects initiated the development of associative learning models such as that proposed by Rescorla and Wagner (1972). Following Kamin's suggestion that the "surprisingness" of an outcome determines the extent that events become associated (Kamin, 1969a, 1969b), Rescorla and Wagner developed a model of Pavlovian conditioning based on the difference between the expected status of an outcome and its actual status (see Allan, 1993, for review). The more unexpected or surprising an outcome, the more conditioning will occur. According to the model, the strength of association between a cue

Table 1.2: Experimental Design by Wagner et al. (1968)

	Learning Phase	Test
Group 1	AB → O (100%)	A ↓
	AC → O (0%)	
Group 2	AB → O (50%)	A ↑
	AC → O (50%)	

and outcome changes as a function of the equation

$$\Delta V = \alpha\beta\left(\lambda - \sum V\right), \tag{1.1}$$

where ΔV represents the change in associative strength of the cue. α is the learning rate parameter that is unique to each cue and represents its salience; it is positive when the cue is present, and zero when it is absent. β is the learning rate parameter associated with the outcome. λ is the upper limit of associative strength that the outcome will support. Finally, $\sum V$ is the sum of associative strengths for all of the cues present on a given trial. Because the outcome can support only a limited amount of associative strength, each cue that is presented must compete to be associated with the outcome.

In Group 1 of Kamin's blocking experiment, the associative strength between Cue A and the outcome quickly approaches asymptote (λ) in the first phase. In the second phase, when the redundant B cue is presented with A, because A has already acquired most of the associative strength available, the sum of associative strength for A and B ($\sum V$) is already near λ. Therefore, the level of surprise ($\lambda - \sum V$) is nearly zero and very little associative strength would accrue to B.[1] At test, Rescorla and Wagner's model predicts very little responding to Cue B.

In the relative validity experiment by Wagner et al. (1968), the individual cues compete to be associated with the outcome just as they did in the blocking paradigm. In Group 1, Cues B and C perfectly predict the presence and absence of the outcome respectively, the sum of the associative strength ($\sum V$) will quickly approach asymptote and little will be learned about Cue A . In Group 2, because B and C predict the outcome only 50% of the time, they require more trials to reach asymptote. As a result, A can accrue more associative strength. Therefore, the animals will respond more at test to A in Group 2 compared to Group 1. Because B quickly approaches λ in Group 1, thereby reducing the associative strength available to A, some

[1]The number of trials in Phase 1, the learning rate parameters for the Cues A and B, and the asymptotic value of the outcome determines the magnitude of $\sum V$. If the associative strength between A and the outcome reaches asymptote, then there is no discrepancy between the actual and expected status of the outcome (i.e., $\lambda - \sum V = 0$) and B will acquire no associative strength. If the associative strength between A and the outcome does not reach asymptote in Phase 1, then there will be a slight discrepancy between λ and $\sum V$ resulting in a small amount of associative strength accrued to Cue B.

would say that Cue B "blocks" conditioning to A (e.g., Baker, Mercier, Vallée-Tourangeau, Frank, & Pan, 1993; Vallée-Tourangeau, Baker, & Mercier, 1994; Baker, Murphy, & Vallée-Tourangeau, 1996; Mehta, 2000). In addition to blocking and relative validity, the competitive learning process has been successful in predicting several counterintuitive phenomena such as overshadowing (Mackintosh, 1975), conditioned inhibition (Williams, 1996), among others (see Siegel & Allan, 1996, for review).

1.2 Blocking Effects in Human Learning

Wasserman (1990) extended the relative validity paradigm to human participants by asking them to judge the efficacy of certain foods in causing an allergic reaction. As in the original relative validity experiment, a common A cue (e.g., Shrimp) was paired with either Cue B (e.g., Strawberries) or C (e.g., Peanuts). The "differential correlation" of the AB and AC compounds with the allergic reaction (O) was varied across five conditions maintaining the original 100-0% and 50-50% manipulation, as well as three intermediate conditions, i.e., 87.5-12.5%, 75-25%, and 62.5-37.5%. Participants' ratings of the redundant A cue increased as a function of the differential correlation between the AB and AC compounds. That is, judgements of A were lowest in the 100-0% condition and highest in the 50-50% condition, and gradually increased among the three intermediate conditions.

Blocking was first demonstrated in humans by Shanks and his colleagues (Dickinson, Shanks, & Evendon, 1984; Dickinson & Shanks, 1985; Shanks, 1985) using a computer game. The experiment followed the two-phase blocking paradigm described earlier. During the first phase of the experiment, participants watched a series of trials where a tank successfully or unsuccessfully traversed a minefield (A $\rightarrow$ O). In the second phase, they were instructed to shoot down the tank. The participants were unable to determine whether the tank was destroyed by a mine or their own gunfire (AB $\rightarrow$ O). During the test phase, they were asked to rate both the influence of their shooting and the effectiveness of the mines in destroying the tank. Those who were exposed to the initial phase of the experiment rated the their own gunfire as being less effective compared to those not exposed to Phase 1. Learning that the minefield was effective in the destruction of the tank seemed to have blocked learning about the influence of their own gunfire in the latter phase of the experiment.

Subsequently, Chapman and Robbins (1990) investigated blocking by presenting participants with information about a fictitious stock market and individual stocks. Over a series of trials, participants were told whether the price of individual stocks increased or not, followed by information about the rise or fall of the entire stock market. The objective was to indicate how predictive each stock was in the fluctuation of the stock market. In the first phase of the experiment, a rise in Stock P (predictive) always resulted in a rise in the market, whereas a rise in Stock N (nonpredictive) resulted in no change in the market. In the second phase, two novel stocks were presented alongside Stocks P and N: Stock B (blocking) was paired with P resulting in a rise in the market; and the pairing of Stock C (control) and N also resulted in a rise in the stock market. During the test phase, participants provided higher predictiveness ratings for C than for B, even though the two stocks were equally predictive. As

in the previous blocking demonstrations, learning about the predictive cue in the first phase blocks learning about the cue it is paired with in the subsequent phase. P blocked the equally predictive association between B and the stock market.

1.2.1 Causal-Model Theory

As researchers started applying the principles of associative learning theories to humans, Waldmann and Holyoak (1992) argued that humans are capable of more sophisticated forms of causal learning than simply reacting to contingencies in their environment. According to Waldmann, people conceptualise the asymmetry of causal relationships. Causes influence effects, but effects do not influence causes. "In addition to using perceived or imagined causes to predict future events, people can use perceived or imagined effects as cues to diagnose their unseen causes." (Waldmann & Holyoak, 1997, p. 125). Our knowledge of causal asymmetry provides us with the capacity to ignore the order that events are presented thereby transforming them into causal-model representations that reflect their asymmetry (Waldmann, 2000). The Rescorla-Wagner model (which embodies the essential and salient characteristics of associative models) neglects the causal status among events by simply encoding their temporal order. Events that occur first are encoded as cues, and subsequent events are encoded as outcomes. It follows from causal-model theory that causes interact and effects do not. That is, we judge one *cause* in light of another, but judge two *effects* independently. According to the Rescorla-Wagner model, *cues* compete and *outcomes* do not. The term *cue-interaction* refers broadly to the relative assessment of two events without reference to the mechanism of interaction. While according to associative models, cues "compete" to be associated with an outcome, according to causal-model theory, people "conditionalise" only among causes. Cue-interaction refers generically to both phenomena.

Chapter 2

Cue-interaction and Judgments of Causality: Contributions of Causal and Associative Processes

> One cannot make causes wiggle in any definite way by manipulating their effects.
> *Hausman (1993)*

2.1 Preface

This chapter is reproduced from Tangen and Allan (submitted). The paper was first submitted on December 18, 2002, and the revision was resubmitted on May 20, 2003 and is currently under review. We wrote this paper in response to the debate between causal-model and associative theorists who have been trying to show that their explanation is the right explanation - and that the other is wrong. We provide evidence that both theoretical camps are correct under different circumstances. Our data highlight the intersection between basic associative processes and our abstract knowledge about causal asymmetry.

2.2 Abstract

In four experiments, the predictions made by causal-model theory and the Rescorla-Wagner model are tested by using a cue-interaction paradigm that measures the relative response to a given event based on the influence or salience of an alternative event. Experiments 1 and 2 uncorrelate two variables that have typically been confounded in the literature (causal order and the number of cues and outcomes) and demonstrate that overall contingency judgments are influenced by the causal structure of the events. Experiment 3 shows that trial-by-trial prediction responses, a second measure of causal assessment, are not influenced by the causal structure of the described events. Experiment 4 revealed that participants became less sensitive to the influence of the causal structure in both their ratings and their predictions as

trials progressed. Thus, two experiments provide evidence for high-level (causal reasoning) processes, and two experiments provide evidence for low-level (associative) processes. We argue that both factors influence causal assessment depending on what is being asked about the events, and participants' experience with those events.

2.3 Introduction

In the past decade, the debate between causal-model and associative learning theorists has centered on whether or not human inferences are sensitive to the causal structure of contingent events (see Waldmann, 2000, for review). While causal models code events in terms of *causes* and *effects*, associative models disregard the causal description of the events, instead coding them solely in terms of their temporal order in which antecedent events are referred to as *cues* and subsequent events as *outcomes*. The disagreement has been with regard to the nature of the processes involved in making causal inferences. According to causal-model theory, expectations of causal structure guide learning about the relevant causal events in a top-down fashion. In contrast, an associative account maintains that causal learning is modelled by the bottom-up acquisition of associative weights guided by simple event pairings. This article examines the extent and circumstances in which these two factors influence causal assessments and describes the conditions under which they operate.

As researchers started applying the principles of associative learning theories to humans (e.g., Shanks & Dickinson, 1987), Waldmann and Holyoak (1992) argued that humans are capable of more sophisticated forms of causal learning than simply reacting to contingencies in their environment. They argued that people conceptualise the asymmetry of causal relationships. Causes influence effects, but effects do not influence causes. "In addition to using perceived or imagined causes to predict future events, people can use perceived or imagined effects as cues to diagnose their unseen causes." (Waldmann & Holyoak, 1997, p. 125). Our knowledge of causal asymmetry provides us with the capacity to ignore the order that events are presented thereby transforming them into causal-model representations that reflect their asymmetry (Waldmann, 2000). The Rescorla-Wagner model (which embodies the essential and salient characteristics of associative models) neglects the causal status among events by simply encoding their temporal order. Events that occur first are encoded as cues, and subsequent events are encoded as outcomes. It follows from causal-model theory that causes interact and effects do not. That is, we judge one *cause* in light of another, but judge two *effects* independently. According to the Rescorla-Wagner model, *cues* compete and *outcomes* do not. The term *cue-interaction* refers broadly to the relative assessment of two events without reference to the mechanism of interaction.

Causal-model and associative theories have often been pitted against one another in the context of cue-interaction paradigms such as blocking (e.g., Waldmann & Holyoak, 1992; Waldmann, 2000), relative cue validity (e.g., Van Hamme, Kao, & Wasserman, 1993; Matute, Arcediano, & Miller, 1996), and overshadowing (e.g., Waldmann, 2001). Of interest in each of these paradigms is the extent to which participants regard one cue in light of another, or consider each cue independently. In the present series of experiments, the one-phase simultaneous

blocking task (Baker et al., 1993) was used to provide a novel test of causal-model theory by means of the conditional ΔP account (Spellman, 1996a, 1996b). According to causal-model theory, when two causes produce one effect, one should consider each cause conditional upon the other because causes interact. When one cause produces two effects, one should consider each effect independent of the other because effects do not interact. The one-phase simultaneous blocking design enables the use of two differentially predictive causes in which the participant is free to conditionalize on one another thereby providing an strong test of the model's predictions. When two causes produce one effect, a conditional ΔP account applied to causal-model theory predicts that participants should rate the influence of each cause in accordance with *conditional* ΔP. When one cause produces two effects participants should rate the influence of the cause on each effect in accordance with *unconditional* ΔP.

In a task involving two cues and a single outcome, one of four cue combinations are possible on a given trial: both cues may be present (AB), one cue may be present and the other absent (A~B or ~AB), or both cues may be absent (~A~B). For each cue combination, the outcome either occurs (O) or does not occur (~O) resulting in eight possible cue-outcome combinations as illustrated in Figure 2.1. Thus, each cue can be expressed in terms of its respective *unconditional* ΔP value defined as:

$$\Delta P_A = P(O|A) - P(O|\sim A) = \frac{a+c}{a+b+c+d} - \frac{e+g}{e+f+g+h} \tag{2.1}$$

$$\Delta P_B = P(O|B) - P(O|\sim B) = \frac{a+e}{a+b+e+f} - \frac{c+g}{c+d+g+h} \tag{2.2}$$

where each equation corresponds to the difference between the proportion of times the outcome occurs given the cue and the proportion of times the outcome occurs not given the cue (Allan, 1980). Alternatively, Cues A and B can be expressed in terms of their respective *conditional* ΔP values defined as:

$$\Delta P_{A|B} = P(O|AB) - P(O|\sim AB) = \frac{a}{a+b} - \frac{e}{e+f} \tag{2.3}$$

$$\Delta P_{A|\sim B} = P(O|A\sim B) - P(O|\sim A\sim B) = \frac{c}{c+d} - \frac{g}{g+h} \tag{2.4}$$

$$\Delta P_{B|A} = P(O|BA) - P(O|\sim BA) = \frac{a}{a+b} - \frac{c}{c+d} \tag{2.5}$$

$$\Delta P_{B|\sim A} = P(O|B\sim A) - P(O|\sim B\sim A) = \frac{e}{e+f} - \frac{g}{g+h} \tag{2.6}$$

The conditional ΔP values in Equations 2.3 - 2.6 allow one to assess the influence of each cue both in the presence and absence of the other cue. For example, to assess the influence of Cue A, Equation 2.3 describes only the cases in which Cue B is present by taking the difference between the proportion of times the outcome occurs given A and the proportion of times the outcome occurs not given A. Moreover, Equation 2.4 describes only the cases in

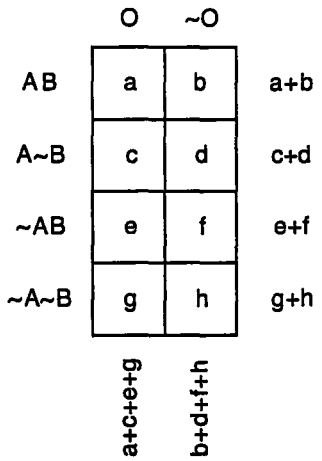

Figure 2.1: Summary 4×2 contingency matrix illustrating each of the possible cause-effect combinations for two cues. Each cell represents the frequency of each event type.

which Cue B is absent by taking the difference between the proportion of times the outcome occurs given A and the proportion of times the outcome occurs not given A.

Therefore, when two causes produce one effect, a conditional ΔP account applied to causal-model theory predicts that, because each cause should be assessed in light of the other, participants should rate the influence of each cause in accordance with conditional ΔP (Equations 2.3 - 2.6). When one cause produces two effects, because each effect should be assessed independently, participants should rate the influence of the cause on each effect in accordance with unconditional ΔP (Equations 2.1 and 2.2). Under these circumstances, with only one cause and two effects, one must rotate the 4×2 contingency matrix shown in Figure 2.1 to form a 2×4 matrix in which the two rows represent the presence and absence of the cause, and the columns represent the four combinations of the two effects. By doing so, it is impossible to calculate the conditional contingencies for A and B defined in Equations 2.3 - 2.6.

Experiments designed to test causal-model theory have typically compared two causal scenarios in which two (or more) causes precede a single effect or in which two (or more) effects precede a single cause thereby confounding causal order (CE vs. EC) and the number of causes and effects (2-1 vs. 1-2) (e.g., Waldmann & Holyoak, 1992; Van Hamme & Wasserman, 1993; Matute et al., 1996; Waldmann, 2000). As illustrated in Figure 2.2, four cause-effect scenarios are possible by crossing the two variables: two cues can be followed by one outcome and be described as two causes producing an effect (2C-1E) or as two effects resulting from a cause (2E-1C), and one cue can be followed by two outcomes and be described as a cause producing two effects (1C-2E) or as an effect resulting from two causes (1E-2C). According to causal-model theory, participants should be sensitive to the interaction between causal order and the number of the causes and effects, which together define the *structure* of the causal relationship (Waldmann & Holyoak, 1992, 1997; Waldmann, 2000, 2001). The model predicts that pairs of causes will interact in the 2C-1E and 1E-2C scenarios (i.e., the negative diagonal

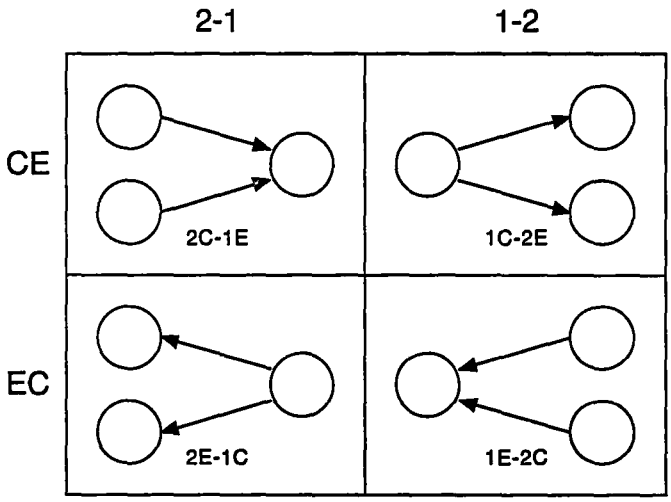

Figure 2.2: Four possible causal scenarios generated by crossing causal order (CE vs. EC) with the number of cues and outcomes (2-1 vs. 1-2).

of Figure 2.2) and predicts that pairs of effects will not interact in the 2E-1C and 1C-2E scenarios (i.e., the positive diagonal of Figure 2.2). In contrast, according to the Rescorla-Wagner model, participants should be sensitive only to the number of the cues and outcomes in which *cues* interact regardless of their causal order. The model therefore predicts that pairs of cues will interact in the 2C-1E and 2E-1C scenarios (i.e., the left column of Figure 2.2) and predicts that pairs of outcomes will not interact in the 1C-2E and 1E-2C scenarios (i.e., the right column of Figure 2.2).

To summarize, a conditional ΔP account applied to causal-model theory predicts that judgments of a pair of differentially predictive *causes* should elicit a cue-interaction effect, while judgments of a pair of differentially diagnostic *effects* should not. In contrast, the Rescorla-Wagner model predicts that judgments of a pair of differentially contingent *cues* should elicit a cue-interaction effect, while judgments of a pair of differentially contingent *outcomes* should not.

2.4 Experiment 1

Experiment 1 was designed to test the predictions made by causal-model theory and the Rescorla-Wagner model by independently manipulating causal order and the number of cues and outcomes. Thus, four causal scenarios were presented to participants using the one-phase simultaneous blocking task described above. Two cues were described either as causes of an effect (2C-1E) or as effects of a cause (2E-1C); or one cue was described either as a cause of two effects (1C-2E) or as an effect of two causes (1E-2C) as shown in Figure 2.2. In each of the four scenarios, the two events that were presented simultaneously were either differentially

Table 2.1: Frequency of events in Experiment 1. The unconditional ΔP values were calculated using Equations 2.1 and 2.2. The conditional ΔP values were calculated using Equations 2.3-2.6.

Trial Type	0.5/0	0.5/1
ABO	9	18
A~BO	9	0
~ABO	3	6
~A~BO	3	0
AB~O	3	0
A~B~O	3	6
~AB~O	9	0
~A~B~O	9	18
# of Trials	48	48
ΔP_A	0.5	0.5
$\Delta P_{A\|B}$	0.5	0
$\Delta P_{A\|\sim B}$	0.5	0
ΔP_B	0	1
$\Delta P_{B\|A}$	0	1
$\Delta P_{B\|\sim A}$	0	1

predictive or diagnostic of the single event. Event A had a moderately positive unconditional ΔP of 0.5, and was paired with B which had an unconditional ΔP of 0 or 1. Causal-model theory predicts that participants will demonstrate a cue-interaction effect in the 2C-1E and 1E-2C scenarios (and not in the other two), and the Rescorla-Wagner model predicts that participants will demonstrate a cue-interaction effect in the 2C-1E and 2E-1C scenarios (and not in the other two).

2.4.1 Method

Participants and Design

Forty-eight undergraduate students at McMaster University participated for course credit. The experiment was designed to test how ratings of a moderately positive contingency varied in the presence of a zero or a perfect contingency as a function of causal order and the number of cues and outcomes. A four-factor mixed design was used with causal order as a between factor with two levels (CE and EC); and the number of cues and outcomes as a within factor with two levels (2-1 and 1-2). Thus, half of the participants were assigned to the CE group and presented with the 2C-1E and 1C-2E scenarios, and half were assigned to the EC group and received the 2E-1C and 1E-2C scenarios. Within each group, the order that the scenarios

were presented was counterbalanced. A third within factor was the contingency of Event B ($\Delta P_B = 0$ and $\Delta P_B = 1$), in which the order of presentation was also counterbalanced. The fourth factor was a within factor representing the number of trials prior to the participants' ratings (32 and 48). Table 2.1 illustrates the trial frequencies obtained by combining an unconditional contingency for A ($\Delta P_A = 0.5$) with one of two unconditional contingencies for B: a zero contingency ($\Delta P_B = 0$) or a perfect contingency ($\Delta P_B = 1$). We use the notation introduced by Baker et al. (1993) to represent the unconditional contingencies of the two events, $\Delta P_A/\Delta P_B$. The designation for the two examples in Table 2.1 are 0.5/0 and 0.5/1, in which the value on the left of the solidus represents ΔP_A and the value on the right represents ΔP_B.

Procedure and Materials

The design and procedure for Experiment 1 were adapted from Mehta (2000). Participants received instructions on a computer screen where they were informed about four strains of bacteria that have been discovered in the mammalian digestive system. In the 2C-1E and 1E-2C scenarios, they were told that scientists were testing whether a pair of chemicals affected the strain's survival, whereas, in the 2E-1C and 1C-2E scenarios, the scientists were testing whether the bacteria affected the production of a pair of chemicals.

Up to four participants at a time performed the experiment on Power Macintosh computers located in separate rooms. The entire experiment was programmed in MetaCard 2.3.1. In the instructions, the four causal scenarios were identified as separate "experiments" designed to test the influence of the chemicals on the bacterial strain, or vice versa. Within each scenario, forty-eight trials were presented in random order according to the frequencies presented in Table 2.1. The addition or production of a chemical was indicated by a computer rendered movie of a colored three-dimensional chemical spinning along its axis, and actual footage of moving bacteria was displayed when the bacterial strain survived or was added. Faded, unmoving greyscale images of the same chemicals and bacteria were displayed to indicate their absence on a given trial. The names of the chemicals and bacteria were displayed only when the events occurred. Each of the movies and images were randomly assigned fictitious names from a set of eight chemicals and four bacteria. Chemical A was always presented on the left hand side of the display, and Chemical B was always presented on the right. The observer initiated a condition by clicking the "Begin" button on the computer screen and initiated each subsequent trial by clicking the "Next Trial" button.

The materials for the four causal scenarios are described as follows:

2C-1E: Participants were instructed that each of the two chemicals would either be added to the bacterial strain or not, resulting in the survival or death of the bacterial strain. They were then presented with a series of trials in which one, both, or neither chemical was added, followed by the survival or death of the bacterial strain.

1C-2E: Participants were instructed that the bacterial strain would either be added to a human digestive environment or not, resulting in the production of each of a pair chemicals or not. They were then presented with a series of trials in which the bacterial strain was either added or not, followed by the production of one, both, or neither chemical.

2E-1C: Participants were instructed that the bacterial strain would either be added to a human digestive environment or not, resulting in the production of each of a pair chemicals or not. They were then presented with a series of trials in which one, both, or neither chemical was produced, followed by the addition of the bacterial strain or not.

1E-2C: Participants were instructed that each of the two chemicals would either be added to the bacterial strain or not, resulting in the survival or death of the bacterial strain. They were then presented with a series of trials in which the bacterial strain survived or not, followed by the addition of one, both, or neither chemical.

After passively viewing a series of thirty-two trials, participants in the 2C-1E and 1E-2C scenarios were asked to rate how strongly each chemical affected the survival of the bacteria, and those in the 2E-1C and 1C-2E scenarios were asked to rate how strongly the bacteria affected the production of each chemical. Ratings were made on a scale ranging from -100 to 100 by moving a horizontal scrollbar with a mouse ranging from -100 at the leftmost position to 100 at the rightmost position, anchored at 0 at the center. After rating A, they were prompted to rate B, followed by another sixteen trials in which they would repeat the rating process. After observing two "experiments" in which ΔP_B was either 0 or 1, a second set of instructions was presented, nearly identical to the first differing only in the number of cues and outcomes as described previously. Again, ΔP_B was either 0 or 1 for the latter two "experiments" comprising a total of four conditions.

Results and Discussion

Mean ratings of Event A after 48 trials are illustrated in Figure 2.3a (error-bars represent standard errors of the means). Ratings for each of the four scenarios are plotted as a function of the two ΔP_B values. According to causal-model theory, when two causes produce a single effect (2C-1E and 1E-2C) ratings of A, which were always moderately positive, should remain moderately positive in the presence of a zero contingency and should be much less positive in the presence of a perfect contingency (tracking the conditional ΔP values in Table 2.1). When two effects result from a single cause (2E-1C and 1C-2E), A should be rated as moderately positive both in the presence of a zero or perfect contingency (tracking the unconditional ΔP values in Table 2.1). According to the Rescorla-Wagner model, cue-interaction should be present only in the 2C-1E and 2E-1C scenarios. The pattern of results presented in Figure 2.3a are consistent with causal-model theory. Only in the 2C-1E and 1E-2C scenarios are ratings of the moderately positive contingency noticeably lower in the presence of a perfect contingency ($\Delta P_B = 1$) than in the presence of a zero contingency ($\Delta P_B = 0$). Although "noticeably lower" here refers to a sizeable negative rating of A, what is relevant is that the *trend* in participants' ratings of A demonstrate conditionalization (see also Spellman, 1996a).

A four-way mixed ANOVA (effects were assessed for significance at the $\alpha = .05$ level), with ratings of A as the dependent variable, revealed significant main effects of contingency for Event B ($\Delta P_B = 0$ vs. $\Delta P_B = 1$), $F(1, 46) = 40.12$, $MSe = 3112.23$, and the number of cues and outcomes (2-1 vs. 1-2), $F(1, 46) = 4.55$, $MSe = 1388.45$. The trial main effect (rating after 32 vs. 48 trials) was not significant, $F(1, 46) = .28$, $MSe = 609.23$, nor did it interact with any of the other factors. The main effect of causal order (CE vs. EC), while not

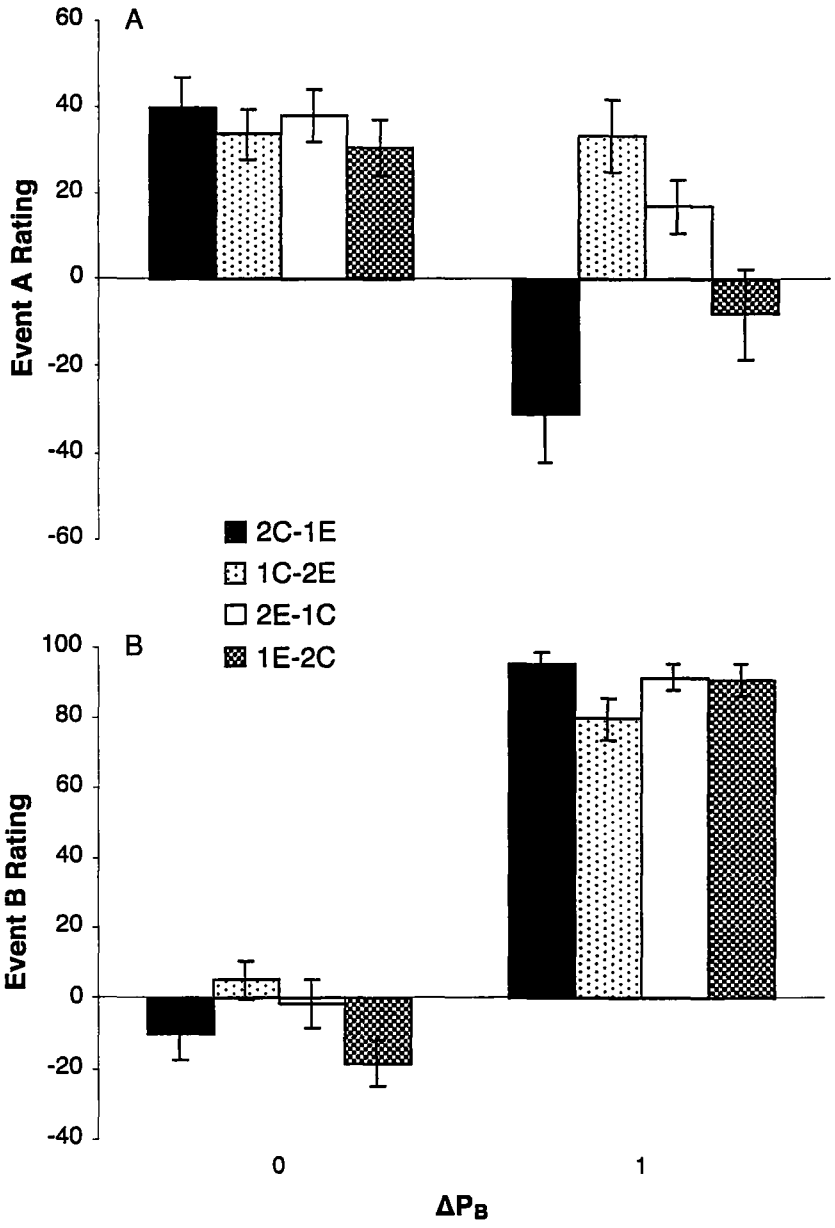

Figure 2.3: Mean ratings in Experiment 1 after 48 trials of Event A (Figure 2.3a) and of Event B (Figure 2.3b). For each event, the ratings are shown as a function of ΔP_B separately for each of the four scenarios. Error-bars represent standard errors of the means.

significant, $F(1, 46) = .04$, $MSe = 2573.04$, did interact with the number of cues and outcomes and the contingency for Event B, $F(1, 46) = 17.78$, $MSe = 2635.78$. This significant three-way interaction was further examined using the Tukey test. When $\Delta P_B = 0$, the ratings were not significantly different among the four scenarios. Moreover, these ratings did not differ from the ratings in the two scenarios in which one cause produced two effects (1C-2E and 2E-1C) when $\Delta P_B = 1$. In contrast, ratings in the two scenarios in which two causes produced one effect (2C-1E and 1E-2C) when $\Delta P_B = 1$ were significantly lower than the other ratings and did not differ from each other.

Mean ratings of Event B are shown in Figure 2.3b. Table 2.1 indicates that for both 0.5/0 and 0.5/1, the conditional probabilities are the same as the unconditional probabilities. Therefore, ratings of B should be the same for the four scenarios, and should be lower for 0.5/0 than for 0.5/1. It is clear from Figure 2.3b that the ratings for B are consistent with causal-model theory. With ratings of Event B as the dependent variable, a four-way ANOVA revealed that the only main effect that was significant was ΔP_B, $F(1, 46) = 628.13$, $MSe = 1394.64$. None of the interactions involving ΔP_B were significant confirming that ratings for a constant ΔP_B did not differ across causal order or number of cue and outcomes. The only other significant outcome was the interaction between causal order and trial, $F(1, 46) = 4.14$, $MSe = 474.94$. The Tukey test revealed that this interaction reflected higher ratings for the CE order than for the EC order after 32 trials but not after 48 trials.

In summary, Experiment 1 resulted in a significant interaction between causal order and the number of cues and outcomes. When two causes resulted in one effect (2C-1E and 1E-2C), participants rated the moderately contingent Cause A as less predictive when it was paired with a perfect predictor ($\Delta P_B = 1$) than when it was paired with a non-predictor ($\Delta P_B = 0$). When one cause resulted in two effects (2E-1C and 1C-2E), participants rated the moderately contingent Effect A as equally diagnostic, both when the effect it had been paired with was perfectly diagnostic ($\Delta P_B = 1$) or was non-diagnostic ($\Delta P_B = 0$). These results indicate that cue-interaction occurs when two causes produce one effect regardless of whether the causes are presented before or after the effect, thus providing clear support for causal-model theory. Participants' overall ratings seem to be sensitive to the causal structure of contingent events.

Both causal-model theory and the Rescorla-Wagner model predict a cue-interaction effect when two causes precede a single effect (2C-1E) and no cue-interaction when one cause precedes two effects (1C-2E). However, only causal-model theory predicts the pattern of results obtained in Experiment 1 in which a cue-interaction effect occurs when one effect precedes two causes (1E-2C) and no cue-interaction occurs when two effects precede one cause (2E-1C). Notice, however, that ratings of A in the presence of a perfect predictor are lower in the 2E-1C scenario compared to the 1C-2E scenario. Similarly, ratings of A in the presence of a perfect predictor are lower in the 2C-1E scenario compared to the 1E-2C scenario. According to causal-model theory, when two effects precede a single cause (2E-1C), there should be no difference between ratings of A when B is perfectly predictive or non-predictive and these ratings should not differ from the 1C-2E scenario. In contrast, when one effect precedes two causes (1E-2C), there *should* be a difference between ratings of A when B is perfectly predictive or non-predictive and these ratings should not differ from the 2C-1E scenario. The data indicate, however, that when the effects come first, the influence of the causal model seems

Table 2.2: Frequency of events in Experiments 2 and 3. Unconditional ΔP values were calculated using Equations 2.1 and 2.2. Conditional ΔP values were calculated using Equations 2.3-2.6.

Trial Type	0.5/0	0.5/0.25	0.5/0.5	0.5/0.75	0.5/1
ABO	6	8	10	11	12
A∼BO	6	4	2	1	0
∼ABO	2	2	2	3	4
∼A∼BO	2	2	2	1	0
AB∼O	2	2	2	1	0
A∼B∼O	2	2	2	3	4
∼AB∼O	6	4	2	1	0
∼A∼B∼O	6	8	10	11	12
# of Trials	32	32	32	32	32
ΔP_A	0.5	0.5	0.5	0.5	0.5
$\Delta P_{A\|B}$	0.5	0.47	0.33	0.17	0
$\Delta P_{A\|\sim B}$	0.5	0.47	0.33	0.17	0
ΔP_B	0	0.25	0.5	0.75	1
$\Delta P_{B\|A}$	0	0.13	0.33	0.67	1
$\Delta P_{B\|\sim A}$	0	0.13	0.33	0.67	1

to lessen, or perhaps, the influence of an associative mechanism may increase. We will revisit this point in the discussion of Experiment 2.

2.5 Experiment 2

The data provided in Experiment 1 indicate that participants' overall ratings are sensitive to the causal structure of the events. Following the suggestion that causal order and the number of cues and outcomes had been confounded in previous investigations of cue-interaction, four causal scenarios were tested in Experiment 1 in which a moderately positive contingency (ΔP_A = 0.5) was paired either with a zero contingency (ΔP_B = 0) or a perfect contingency (ΔP_B = 1). Experiment 2 was designed to replicate the results from Experiment 1 and to generalize from the extreme contingencies used to less extreme values by including three intermediate ΔP_B values (0.25, 0.5, 0.75). The ΔP_B values chosen for the three intermediate contingency pairs were selected to best contrast the predictions made by the Rescorla-Wagner model and causal-model theory through participants' ratings of A, and were not chosen for their intrinsic value. To clarify, several different frequencies can be selected to fill the eight cells of the 4×2 matrix each resulting in various combinations of unconditional and conditional ΔP values. The frequencies for Experiment 2 (shown in Table 2.2) were selected to produce a descending

pattern of conditional ΔP_A values while maintaining identical unconditional ΔP_A values. As well, they were selected so the unconditional and conditional ΔP_B values would be as closely matched as possible. The ΔP_B values were therefore selected only for their influence on the conditional ΔP_A values. The frequencies were also selected so the respective conditional contingencies for A and B would be identical where $\Delta P_{A|B} = \Delta P_{A|\sim B}$ and $\Delta P_{B|A} = \Delta P_{B|\sim A}$ resulting in the symmetry observed in the two columns of the 4×2 contingency matrix for each of the five conditions (see Spellman, 1996b, Property 4).

In addition, Experiment 2 was designed to independently test each of the four causal scenarios. In Experiment 1, half of the participants were presented with both the 2C-1E and 1C-2E scenarios and the other half were presented with the 2E-1C and 1E-2C scenarios. In Experiment 2, however, each group was presented with only one causal scenario (2C-1E, 2E-1C, 1C-2E, or 1E-2E).

2.5.1 Method

Participants and Design

Sixty undergraduate students at McMaster University participated for course credit. The experiment was designed as a replication of Experiment 1 using five contingency pairs rather than two, casual scenario as a between factor, and a total of 32 rather than 48 trials with a single overall rating. The 60 participants were randomly assigned to one of the four causal scenarios (i.e., 2C-1E, 1C-2E, 2E-1C, or 1E-2C). Within each group, the presentation order of the five ΔP_B values was randomized. Table 2.2 illustrates the trial frequencies obtained by combining $\Delta P_A = 0.5$ with each of the five ΔP_B values.

Procedure and Materials

The procedure and materials in Experiment 2 were very similar to those in Experiment 1. The difference was in the total number trials and the number of ΔP_B values. Participants were presented with 32 trials before rating Events A and B, where they would repeat the process after observing each of the five "experiments". Two more fictitious chemicals and one more bacterial strain was added among those to be presented.

Results and Discussion

Mean ratings of Event A are illustrated in Figure 2.4a. Ratings for each of the four causal scenarios are plotted as a function of the five ΔP_B values. According to causal-model theory, ratings of A in the 2C-1E and 1E-2C scenarios should track the pattern of conditional ΔP_A values presented in Table 2.2. The conditional ΔP_A values decrease as ΔP_B increases, and therefore ratings of A should also decrease. Causal-model theory also predicts that the ratings of A in the 1C-2E and 2E-1C scenarios should track the pattern of unconditional ΔP_A values presented in Table 2.2. The unconditional ΔP_A values are constant, and therefore the ratings of A should not change across the five ΔP_B values. The Rescorla-Wagner model makes similar predictions but for different scenarios: ratings of A should be a decreasing function of ΔP_B

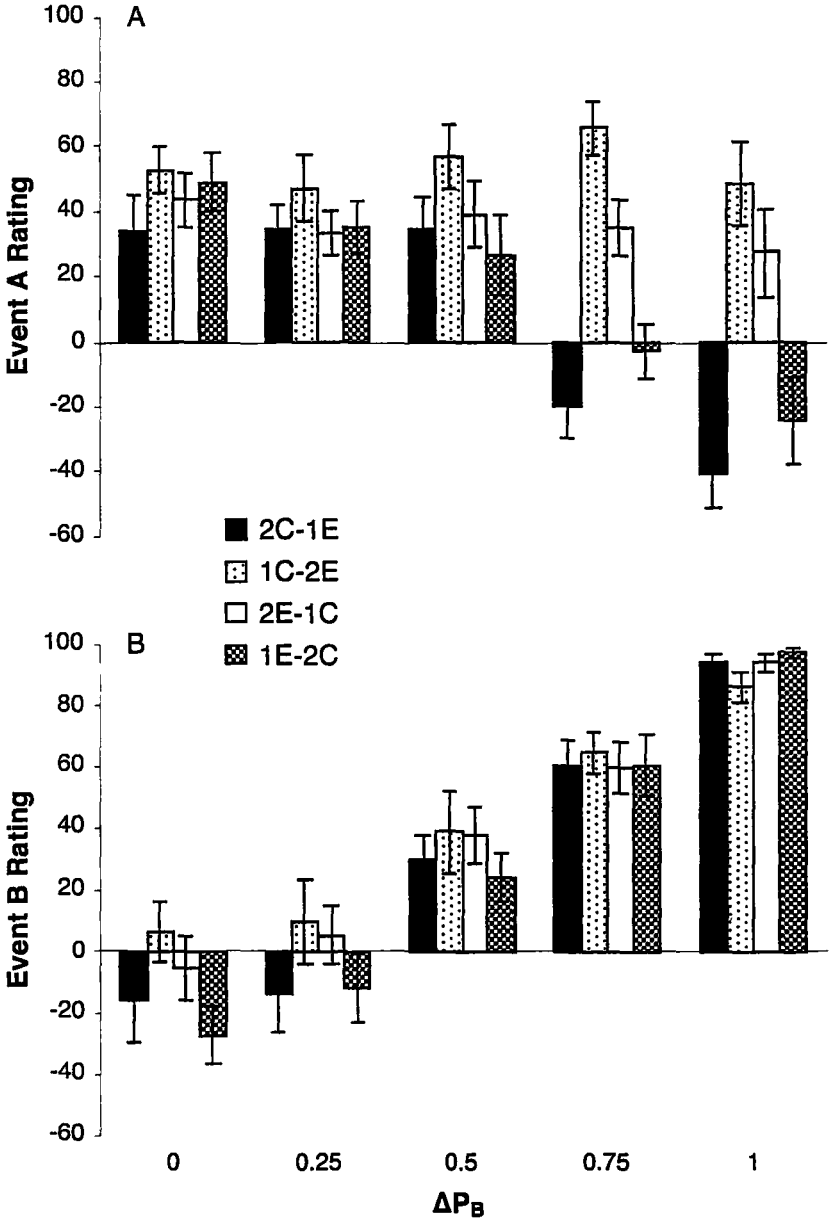

Figure 2.4: Mean ratings in Experiment 2 after 32 trials of Event A (Figure 2.4a) and of Event B (Figure 2.4b). For each event, the ratings are shown as a function of ΔP_B (0, 0.25, 0.5, 0.75, 1) separately for each of the four scenarios. Error-bars represent standard errors of the means.

for the 2C-1E and 2E-1C scenarios, and should be independent of ΔP_B for the 1C-2E and 1E-2C scenarios. As we noted earlier, the predictions for both models are ordinal. Thus, we are examining not only the presence or absence of cue-interaction, but also the ordinal level of cue-interaction among the four causal scenarios.

The ratings of A appear to support the predictions made by causal-model theory. In the 2C-1E and 1E-2C scenarios, ratings of A decline as ΔP_B increases, tracking the pattern of conditional ΔP_A values presented in Table 2.2. In the 2E-1C and 1C-2E scenarios, ratings of A remain relatively constant regardless of the contingency for Event B, tracking the pattern of unconditional ΔP_A values presented in Table 2.2.

With four causal scenarios (2C-1E, 1C-2E, 2E-1C, 1E-2C) as a between factor and five ΔP_B values (0, 0.25, 0.5, 0.75, 1) as a within factor, a mixed ANOVA was conducted on the ratings of A. As expected, the analysis revealed main effects of causal scenario, $F(3, 56) = 14.83$, $MSe = 2098.34$, and ΔP_B, $F(4, 224) = 13.32$, $MSe = 1348.32$, as well as a significant interaction between them, $F(12, 224) = 3.98$, $MSe = 1348.32$. The Tukey test was used to examine this significant interaction to see whether the results replicated those found in Experiment 1. The ratings of A for the two ΔP_B values used in Experiment 1 ($\Delta P_B = 0$ and $\Delta P_B = 1$) were compared and ratings were not significantly different among the four causal scenarios when $\Delta P_B = 0$. Also, these ratings did not differ from the ratings when $\Delta P_B = 1$ if one cause produced two effects (1C-2E and 2E-1C). In contrast, when two causes produced a single effect (2C-1E and 1E-2C) and $\Delta P_B = 1$, ratings were significantly lower than the other ratings and did not differ from each other. Thus, ratings of A in Experiment 2 provide a replication of the Experiment 1 results. Cue-interaction occurs when two causes result in one effect, regardless of whether the causes precede or follow the effect, and cue-interaction does not occur when a single cause results in two effects, regardless of their causal order.

According to causal-model theory, ratings of A should decrease as ΔP_B increases when two causes produce one effect (2C-1E and 1E-2C), and should remain constant when one cause produces two effects (1C-2E and 2E-1C). A linear trend analysis was conducted on the A ratings, separately for each scenario, across the five ΔP_B values[1]. As predicted by causal-model theory, the linear trend was significant for the 2C-1E, $F(1, 56) = 40.10$, and 1E-2C, $F(1, 56) = 29.60$, scenarios, and was not significant for the 1C-2E, $F(1, 56) = .09$, and 2E-1C, $F(1, 56) = .74$, scenarios ($MSe = 1751.69$ for each comparison).

Mean ratings of Event B are illustrated in Figure 2.4b. The ratings of B clearly increase with ΔP_B. Table 2.2 indicates that for 0.5/0.25, 0.5/0.5, and 0.5/0.75, the conditional values of ΔP_B are less than the unconditional values. Thus, according to causal-model theory, ratings of B if two causes produce one effect (2C-1E and 1E-2C) should be less than if one cause produces two effects (1C-2E and 2E-1C). While the data tend in that direction, the statistical analysis indicated that the scenario effect was not significant. With ratings of B as the dependent measure, a mixed ANOVA revealed only a significant main effect of ΔP_B, $F(4, 224) = 107.84$, $MSe = 1047.12$. The main effect of causal scenario was not significant, $F(3,$

[1]We are interested in whether there is a significant linear trend among the A ratings across the five levels of ΔP_B, tracking the conditional ΔP values for Event A. The interval between the levels of the independent variable are unequal (i.e., 0.5, 0.47, 0.33, 0.17, 0) whereby the following coefficients were derived: 21, 17, 4, -13, -29 (see Howell, 1997, for derivation).

56) = 1.16, MSe = 2296.1, nor was the interaction between causal scenario and ΔP_B, $F(12, 224)$ = .89, MSe = 1047.13. To evaluate whether the absence of a significant scenario effect was attributable to the cases where the conditional and unconditional values of ΔP_B were the same (0.5/0 and 0.5/1), the ANOVA was conducted on the three other pairings (0.5/0.25, 0.5/0.5, and 0.5/.75). Again, only the main effect of ΔP_B was significant, $F(2, 112) = 50.51$, MSe = 1223.04.

In summary, the ratings of A in Experiment 2 provide a direct replication of the ratings in Experiment 1 and generalize the results to less extreme ΔP_B values. When two causes produce one effect, participants rated the moderately positive cause as less predictive when it was paired with a strong predictor than when it was paired with a weak predictor. When a single cause produced two effects, participants rated the moderately positive effect as equally diagnostic regardless of the diagnosticity of the effect that it was paired with. This interaction between causal order and the number of cues and outcomes is consistent with the predictions of causal-model theory. Although not statistically significant, the ratings of B were also consistent with causal-model theory. It must be emphasized that the B ratings do not provide a strong assessment of the models because the ΔP_B values were selected only for their influence on the conditional ΔP_A values.

As in Experiment 1, the causal-model effect is not as strong when the effect(s) precede the cause(s). In Experiment 2, we see that ratings of A are consistently lower in the 2E-1C scenario compared to the 1C-2E scenario. Similarly, ratings of A tend to be lower in the 2C-1E scenario compared to the 1E-2C scenario. Again, while the differences are not significant, when the effect(s) precede the cause(s), participants' ratings seem be influenced less by the causal description of the events and more by their associative strength. Although the data from Experiments 1 and 2 provide conclusive evidence that participants' judgments are driven primarily by the structure of the causal relationship, we will demonstrate the significant role of associative processes in the following two experiments.

2.6 Experiment 3

A conditional ΔP account applied to causal-model theory allows one to generate dichotomous predictions in which cue-interaction should occur or not (as has been done in previous investigations), but in addition, it allows for ordinal predictions where the relative effectiveness of each event determines the degree to which they interact. The data from Experiments 1 and 2 provide solid evidence for the influence of causal expectation on human inference. In Experiment 3 we demonstrate that these high level processes may not occur independently of basic, low-level (associative) processes by exploring a different measure of causal assessment.

In Experiments 1 and 2, participants passively viewed a series of trials before providing an overall rating of the relationship between the events. Our methodology differs from that reported by others (e.g., Shanks & López, 1996; Price & Yates, 1995; Cobos, López, Caño, Almaraz, & Shanks, 2002; Waldmann & Holyoak, 1992) who required participants to predict the outcome of each trial, and were provided corrective feedback on their prediction. For example, on each trial, participants in Experiment 1 of Waldmann and Holyoak (1992) would

see descriptions of people on a computer screen and were to use those descriptions to predict whether they thought a person had the described disease (by pressing a *Yes* key) or did not have the disease (by pressing a *No* key). After indicating their response, they received *Correct* or *Incorrect* as feedback. If participants are presented with four types of event combinations (AB, A∼B, ∼AB, ∼A∼B) and are asked to predict the outcome of each trial (Yes, No), then a 4×2 matrix, like the one presented in Figure 2.1, can be constructed where the columns represent the two prediction responses (Yes, No) rather than the actual outcomes. These predictions can then be used as an indirect measure of their conditional ΔP estimates (López, Shanks, Almaraz, & Fernandez, 1998; Tangen & Allan, 2003).

We have shown in Experiments 1 and 2 that participants demonstrate a sensitivity to the structure of causal relationships which is consistent with the predictions made by causal-model theory. To further investigate participants' sensitivity to causal structure, we required participants in Experiment 3 to predict the outcome of each trial in addition to providing an overall rating of the relationship between the events. Thus, we obtained both a measure of causal assessment derived from prediction responses, as well as explicit overall judgments between the events to determine whether the two measures were congruent as we varied the structure of the causal relationship.

Among the four causal scenarios described earlier (2C-1E, 1C-2E, 2E-1C, and 1E-2C), the results from Experiments 1 and 2 reveal that neither causal order or the number of cues and outcomes were significant factors independently. Instead, the important variable was the interaction between the two factors, i.e., the structure of the causal relationship. Therefore, to avoid the potential confound of the number of predictions participants were making on each trial, we eliminated the right hand column of Figure 2.2 and presented them with only two cues and one outcome (2C-1E and 2E-1C). Each group was shown identical stimuli, but the causal description of the stimuli differed between the two groups. According to causal-model theory, judgments should vary depending on whether the events are described as two causes resulting in an effect, or as two effects resulting from one cause. In contrast, the Rescorla-Wagner model does not make a distinction between the causal description of the events, and codes the two scenarios identically as two cues followed by one outcome. On each trial, a participant was presented with one of four event combinations (AB, A∼B, ∼AB, ∼A∼B) and then predicted whether the effect/cause occurred given the information from the preceding pair of events and from previous trials. Corrective feedback (Correct, Incorrect) was provided immediately after making their decision. After 32 trials, they were asked to provide an overall rating of the relationship between the events as in the previous experiments. The same five contingency pairs were used as in Experiment 2.

2.6.1 Method

Participants and Design

Thirty undergraduate students at McMaster University took part in this experiment for course credit. The design of Experiment 3 was identical to that used in Experiment 2 except the 1C-2E and 1E-2C causal scenarios were eliminated, and participants were asked to predict the outcome of each trial and were provided feedback on their decision. The frequency of events

in Experiment 3 are shown in Table 2.2.

Procedure and Materials

The same procedure and materials as Experiment 2 were used with the addition of predictions on each trial. Participants were presented with two cues consisting of the presence or absence of two chemicals (2C or 2E) and were then asked to predict whether they thought the bacterial strain survived/was added or not by clicking one of two buttons on the computer screen. Once they made their selection, they were presented with the outcome (1E or 1C) along with *Correct* or *Incorrect* as feedback. The prediction responses for each event combination were recorded and used to calculate estimated conditional ΔP values by counting the number of *Yes* and *No* responses for each event combination (AB, A∼B, ∼AB, ∼A∼B) after 16, 32, 48, and 64 trials and substituting these frequencies into Equations 2.3 - 2.6.

Results

In this experiment there were two dependent measures, ratings and predictions.

Ratings Figures 2.5a and 2.5b depict the mean ratings for Cues A and B respectively. The pattern of results for both cues is similar to that observed in Experiment 2. Ratings of A, in the 2C-1E scenario, decline as ΔP_B increases, tracking the pattern of conditional ΔP_A values presented in Table 2.2. Ratings of A, in the 2E-1C scenario, remain relatively constant as ΔP_B increases, tracking the pattern of unconditional ΔP_A values presented in Table 2.2. With ratings of A as the dependent variable, a mixed ANOVA, with causal scenario (2C-1E vs. 2E-1C) as a between factor and ΔP_B (0, 0.25, 0.5, 0.75, 1) as a within factor, revealed significant main effects for scenario, $F(1, 28) = 8.03$, $MSe = 4048.77$, and ΔP_B, $F(4, 112) = 5.21$, $MSe = 1845.58$, as well as a significant interaction, $F(4, 112) = 2.95$, $MSe = 1845.58$. As in Experiment 2, the linear trend was significant for the 2C-1E scenario, $F(1, 28) = 23.80$, but not for the 2E-1C scenario, $F(1, 28) = .66$ ($MSe = 2338.93$ for both comparisons).

Figure 2.5b indicates that the ratings of B increase with ΔP_B and do not appear to depend on causal scenario. With ratings of B as the dependent measure, a mixed ANOVA revealed a significant main effect of ΔP_B, $F(4, 112) = 18.59$, $MSe = 2069.06$. The main effect of causal scenario was not significant, $F(1, 28) = .001$, $MSe = 3001.04$, nor was the interaction between contingency and causal scenario, $F(4, 112) = 1.26$, $MSe = 2069.06$.

Predictions Figures 2.6a and 2.6b plot the estimated ΔP values for Cue A conditional on the presence and absence of B respectively. It is clear from these two figures that the participants' prediction responses are at variance with their ratings. A comparison of the two figures also indicates that the estimates of $\Delta P_{A|B}$ are different from the estimates of $\Delta P_{A|\sim B}$. A 2 (scenario: 2C-1E, 2E-1C) × 2 (Cue B status: present, absent) × 5 (ΔP_B: 0, 0.25, 0.5, 0.75, 1) mixed ANOVA on the estimated conditional ΔP values for A confirms these observations. The main effect of causal scenario was not significant, $F(1, 28) = 1.25$, $MSe = 1201.76$, nor did it interact with ΔP_B, $F(4, 112) = 1.00$, $MSe = 1088.09$, Cue B status, $F(1,$

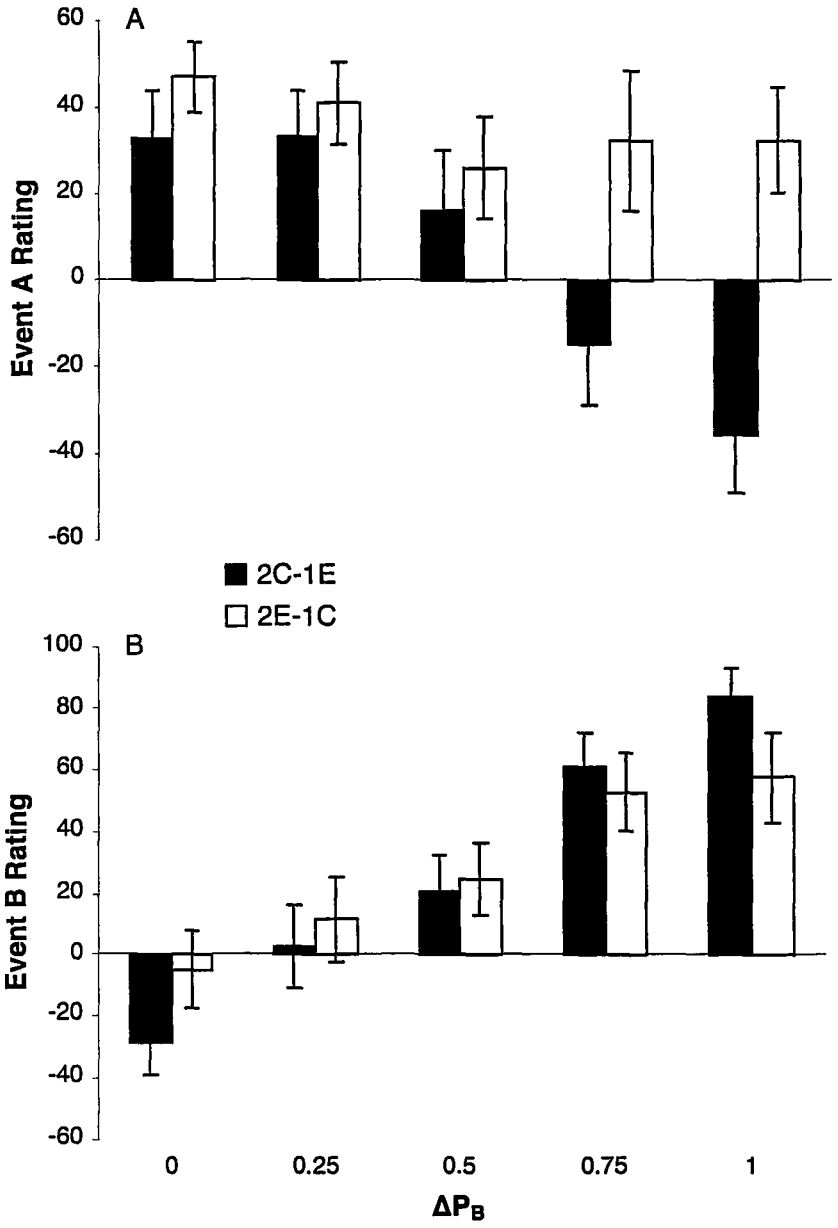

Figure 2.5: Mean ratings in Experiment 3 after 32 trials of Cue A (Figure 2.5a) and of Cue B (Figure 2.5b). For each cue, the ratings are shown as a function of ΔP_B (0, 0.25, 0.5, 0.75, 1) separately for each of the two scenarios. Error-bars represent standard errors of the means.

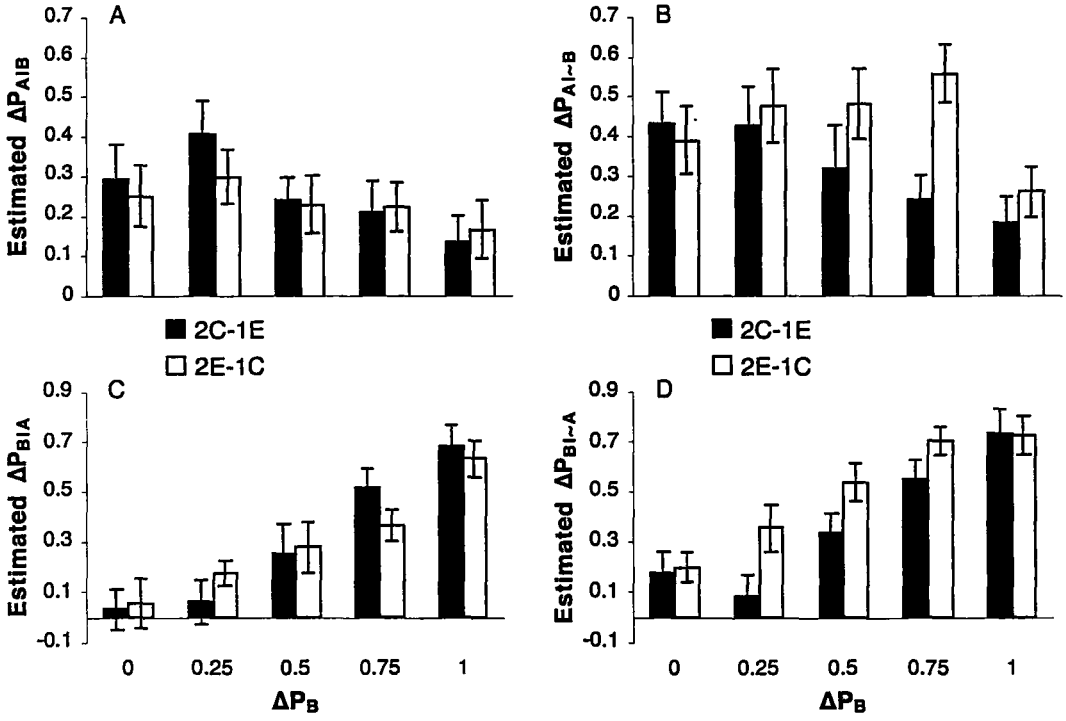

Figure 2.6: Mean estimated conditional ΔP values in Experiment 3 for Cues A and B as a function of ΔP_B separately for the two causal scenarios. $est\Delta P_{A|B}$ is shown in Figure 2.6a, $est\Delta P_{A|\sim B}$ is shown in Figure 2.6b, $est\Delta P_{B|A}$ is shown in Figure 2.6c, and $est\Delta P_{B|\sim A}$ is shown in Figure 2.6d. Error-bars represent standard errors of the means.

28) = 2.54, MSe = 1357.18, or with both, $F(4, 112) = .9$, MSe = 603.56. The main effect of ΔP_B was significant, $F(4, 112) = 3.46$, MSe = 1088.09. The main effect of Cue B status was also significant, $F(1, 28) = 9.91$, MSe = 1357.18, indicating that estimated conditional ΔP for A was lower when B was present ($est\Delta P_{A|B} = .24$) than when it was absent ($est\Delta P_{A|\sim B} = .38$). The interaction between ΔP_B and Cue B status was not significant, $F(4, 112) = .54$, MSe = 603.56.

Figures 2.6c and 2.6d illustrate the estimated ΔP values for B conditional on the presence and absence of A respectively. A 2 (scenario: 2C-1E, 2E-1C) × 2 (Cue A status: present, absent) × 5 (ΔP_B: 0, 0.25, 0.5, 0.75, 1) mixed ANOVA on the estimated conditional ΔP values for B revealed a main effect of ΔP_B, $F(4, 112) = 31.26$, MSe = 1137.41. The main effect of Cue A status was also significant, $F(1, 28) = 10.06$, MSe = 1349.84, indicating that estimated conditional ΔP for B was lower when A was present ($est\Delta P_{B|A} = .31$) than when it was absent ($est\Delta P_{B|\sim A} = .44$). No other effects or interactions reached significance.

Discussion

The results from Experiment 3 provide a direct replication of the rating data obtained in Experiments 1 and 2. Participants rated identical contingencies quite differently depending on whether the events had been described as causes or effects. In the 2C-1E scenario, participants gave lower ratings to the moderately predictive Cause A when it was paired with a highly predictive Cause B than when it was paired with a less predictive Cause B, indicating that causes interact. In contrast, in the 2E-1C scenario, the ratings of Effect A did not depend on the contingency of Effect B, indicating that effects do not interact.

In contrast to the ratings, a causal scenario effect was not seen with the prediction responses. For both 2C-1E and 2E-1C, the estimated conditional ΔP values for A decreased as unconditional ΔP_B increased, indicating that cue-interaction occurred in both scenarios. There appears to be a dissociation between the ratings and the prediction responses. Table 2.2 shows that for each cue, the two conditional ΔP values were always the same. That is, $\Delta P_{A|B} = \Delta P_{A|\sim B}$ and $\Delta P_{B|A} = \Delta P_{B|\sim A}$. This was not the case, however, for the estimates based on the participants predictions, where $est\Delta P_{A|B} < est\Delta P_{A|\sim B}$ and $est\Delta P_{B|A} < est\Delta P_{B|\sim A}$. That is, the estimated conditional ΔP value was smaller when the cue conditionalized upon was present than when it was absent. This pattern of results was also found by Tangen and Allan (2003).

In summary, identical stimuli were presented to participants that were described either as two causes of an effect (2C-1E) or as two effects of a cause (2E-1C). Participants' overall judgments of these relationships varied systematically depending on their causal labels. In addition to making an overall judgment of the relationship, they were asked to make a prediction as to the whether the outcome would occur or not on each trial. Their prediction responses did not vary according to the causal description of the events.

We have revealed a dissociation between two means of assessing judgments of causality. Trial-by-trial prediction responses require participants to estimate the presence or absence of the outcome. The results suggest that participants manage this task by simply basing their judgment on the current level of associative strength, identifying cues as generic events

without any deeper recognition of their causal status. Overall ratings, on the other hand, require participants to not only consider the status of a single outcome, but also take into account the causal relationship among the events presented.

Thus, it seems that participants can report either the current level of associative strength in their predictions by basing their causal assessments on the number of cues and outcomes rather than on the causal structure of the events; or their assessments can reflect the causal status of the events by considering how they are structured. It depends on the nature of the question being asked.

2.7 Experiment 4

The results from Experiment 3 reveal that participants were sensitive to the causal description of the cues and outcome in rating the overall relationship, but the effect was absent in their trial-by-trial predictions. Experiment 4 was designed to further investigate this dissociation between ratings and prediction responses by increasing the total number of trials in each condition from 32 to 64, and having participants provide an overall rating after 16, 32, 48, and 64 trials. By increasing the number of ratings, we can compare each measure across trials as a function of causal scenario. Perhaps a greater number of trials would result in a greater sensitivity to the associative processes at work and less sensitivity to the causal description of the events. Increasing the total number of trials resulted in the elimination of the 0.5/0.5 contingency pair to maintain a one-hour experimental session.

2.7.1 Method

Participants and Design

Forty undergraduate students at McMaster University took part in this experiment for course credit. The design of Experiment 4 was similar to Experiment 3, but the total number of trials was increased to 64, participants were asked to rate each cue after 16, 32, 48, and 64 trials, and the 0.5/0.5 contingency pair was eliminated. The event frequencies in Experiment 4 are shown in Table 2.3.

Procedure and Materials

The procedure and materials for Experiment 4 were similar to those used in Experiment 3 apart from the total number of trials presented and the number of ratings provided by participants. Four contingency pairs were presented to participants as separate "experiments".

Results

As in Experiment 3, there were two dependent measures, ratings and predictions.

Table 2.3: Frequency of events in Experiment 4. Unconditional ΔP values were calculated using Equations 2.1 and 2.2. Conditional ΔP values were calculated using Equations 2.3-2.6.

Trial Type	0.5/0	0.5/0.25	0.5/0.75	0.5/1
ABO	12	16	22	24
A∼BO	12	8	2	0
∼ABO	4	4	6	8
∼A∼BO	4	4	2	0
AB∼O	4	4	2	0
A∼B∼O	4	4	6	8
∼AB∼O	12	8	2	0
∼A∼B∼O	12	16	22	24
# of Trials	64	64	64	64
ΔP_A	0.5	0.5	0.5	0.5
$\Delta P_{A\|B}$	0.5	0.47	0.17	0
$\Delta P_{A\|\sim B}$	0.5	0.47	0.17	0
ΔP_B	0	0.25	0.75	1
$\Delta P_{B\|A}$	0	0.13	0.67	1
$\Delta P_{B\|\sim A}$	0	0.13	0.67	1

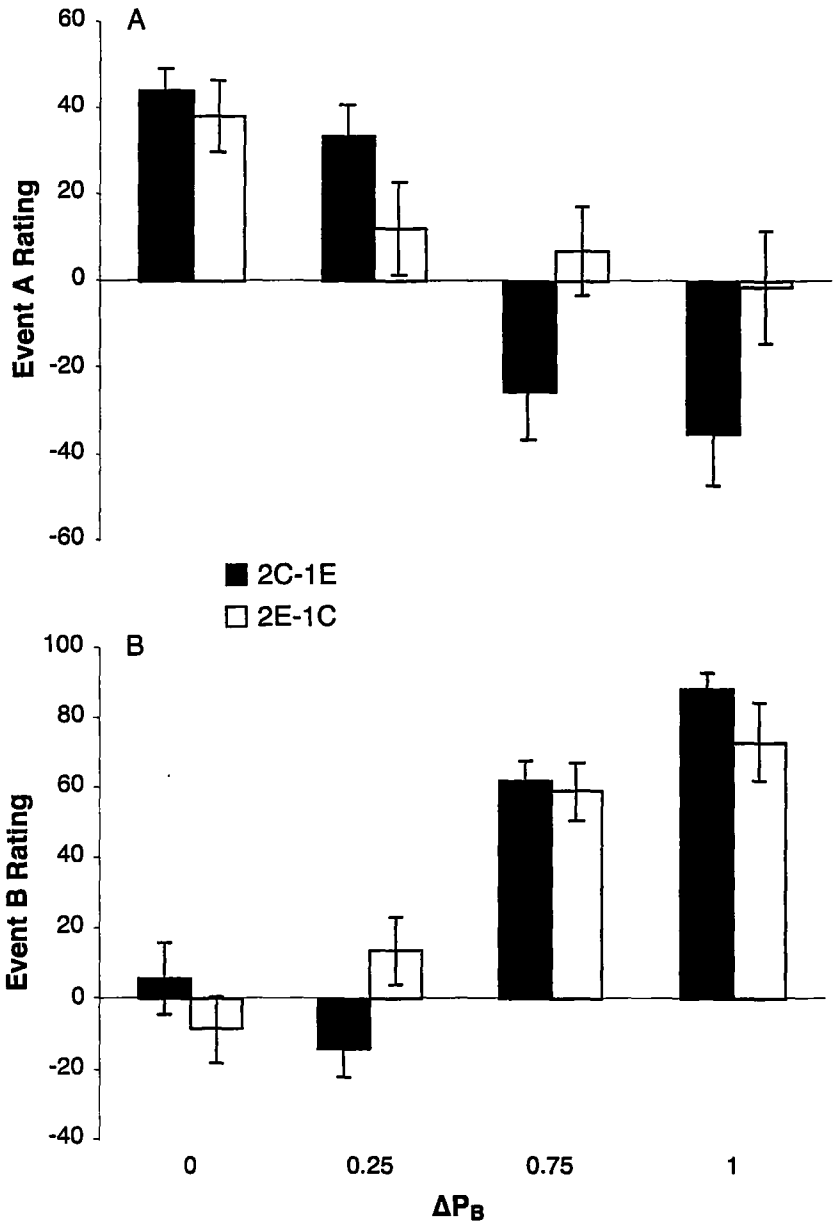

Figure 2.7: Mean ratings in Experiment 4 after 64 trials of Event A (Figure 2.7a) and of Event B (Figure 2.7b). For each cue, the ratings are shown as a function of ΔP_B (0, 0.25, 0.75, 1) separately for each of the two conditions. Error-bars represent standard errors of the means.

Table 2.4: Experiment 4 overall ratings and estimated conditional ΔP values for Cue A after 32, 48, and 64 trials.

		2C-1E							
		0.5/0		0.5/0.25		0.5/0.75		0.5/1	
32	Rating	37.7	(7.6)	14.8	(8.4)	30.7	(13.2)	-38	(11)
	$est\Delta P_{A\|B}$	0.41	(0.07)	0.27	(0.06)	0.17	(0.07)	0.18	(0.07)
	$est\Delta P_{A\|\sim B}$	0.44	(0.1)	0.52	(0.06)	0.31	(0.08)	0.24	(0.06)
48	Rating	29.1	(9.6)	24	(8.6)	-10.8	(10.7)	-32.9	(12.3)
	$est\Delta P_{A\|B}$	0.49	(0.06)	0.32	(0.07)	0.16	(0.06)	0.1	(0.04)
	$est\Delta P_{A\|\sim B}$	0.47	(0.09)	0.51	(0.07)	0.26	(0.05)	0.11	(0.04)
64	Rating	44	(5.1)	33.3	(7.1)	-25.4	(11.1)	-35.3	(11.9)
	$est\Delta P_{A\|B}$	0.51	(0.06)	0.35	(0.07)	0.17	(0.05)	0.08	(0.03)
	$est\Delta P_{A\|\sim B}$	0.51	(0.07)	0.47	(0.06)	0.23	(0.04)	0.08	(0.02)
		2E-1C							
		0.5/0		0.5/0.25		0.5/0.75		0.5/1	
32	Rating	33.6	(9.5)	26	(10.3)	5.7	(11.7)	11.3	(14)
	$est\Delta P_{A\|B}$	0.46	(0.06)	0.38	(0.07)	0.17	(0.09)	0.11	(0.05)
	$est\Delta P_{A\|\sim B}$	0.55	(0.08)	0.56	(0.08)	0.35	(0.07)	0.3	(0.08)
48	Rating	32.4	(8.8)	36.7	(9)	-12.4	(10.1)	-3.15	(13)
	$est\Delta P_{A\|B}$	0.47	(0.06)	0.45	(0.06)	0.14	(0.07)	0.09	(0.04)
	$est\Delta P_{A\|\sim B}$	0.6	(0.07)	0.57	(0.06)	0.3	(0.06)	0.19	(0.05)
64	Rating	37.7	(8.4)	12.1	(10.7)	6.8	(10.3)	-1.5	(12.9)
	$est\Delta P_{A\|B}$	0.5	(0.05)	0.46	(0.06)	0.15	(0.06)	0.08	(0.04)
	$est\Delta P_{A\|\sim B}$	0.58	(0.07)	0.59	(0.06)	0.31	(0.05)	0.17	(0.04)

Note. Standard errors of the means are given in parentheses.

Ratings Figures 2.7a and 2.7b depict the mean ratings after 64 trials for Cues A and B respectively, and Table 2.4 depicts the mean and standard error of the ratings for Cue A after 32, 48, and 64 trials. The ratings and estimated ΔP values after 16 trials are not reported as participants' prediction responses of the randomly presented events occasionally resulted in 4×2 matrices with row frequencies of zero. The pattern of results after 32 trials is similar to that of Experiments 1, 2, and 3. In the 2C-1E scenario, ratings of A roughly approximate the conditional ΔP values presented in Table 2.3, whereas in the 2E-1C scenario the ratings are consistent with the unconditional ΔP values. After 48 and 64 trials, however, a different pattern of results emerges. As illustrated in Figure 2.7a, ratings of A decline as ΔP_B increases, regardless of the causal scenario. The effect of the causal-model seems to have dissipated over trials, and cue-interaction occurs for both scenarios. A 2 (scenario: 2C-1E, 2E-1C) × 4 (ΔP_B: 0, 0.25, 0.75, 1) × 3 (trial: 32, 48, 64) mixed ANOVA on the ratings of A revealed only a significant main effect for ΔP_B, $F(3, 114) = 21.27$, $MSe = 3189.24$, which contributed to significant interactions with scenario, $F(3, 114) = 3.33$, $MSe = 3189.24$, trial, $F(6, 228) = 2.60$, $MSe = 1397.89$, and a three-way interaction with trial and scenario, $F(6, 228) = 3.13$, $MSe = 1397.89$. A linear trend analysis[2] was conducted on the A ratings, separately for each scenario after 32, 48, and 64 trials. For the 2C-1E scenario, the linear trend was significant after 32 trials, $F(1, 38) = 13.59$, $MSe = 2715.31$, 48 trials, $F(1, 38) = 19.22$, $MSe = 2703.87$, and 64 trials, $F(1, 38) = 37.95$, $MSe = 2482.45$. For the 2E-1C scenario, the linear trend was not significant after 32 trials, $F(1, 38) = 2.87$, $MSe = 2715.31$, but was significant after 48 trials, $F(1, 38) = 11.05$, $MSe = 2703.87$, and 64 trials, $F(1, 38) = 4.70$, $MSe = 2482.45$. Thus, by 48 trials, cue-interaction is seen in both scenarios.

Figure 2.7b presents the mean and standard error of the ratings for Cue B after 64 trials, and Table 2.5 presents the mean and standard error of the ratings for Cue B after 32, 48, and 64 trials. Ratings of B seem fairly typical of the results obtained in Experiments 1-3. Mean ratings increase for both scenarios as a function of ΔP_B. A 2 (scenario: 2C-1E, 2E-1C) × 4 (ΔP_B: 0, 0.25, 0.75, 1) × 3 (trial: 32, 48, 64) mixed ANOVA on the ratings of B confirms this observation. The only significant main effect was for ΔP_B, $F(3, 114) = 77.99$, $MSe = 2760.85$. The only other significant effect was a three-way interaction between ΔP_B, trial, and scenario, $F(6, 228) = 3.55$, $MSe = 747.98$, resulting primarily from an exceptionally low mean rating in the 2C-1E scenario, 0.5/0.25 condition, after 64 trials.

Predictions Figures 2.8a and 2.8b plot the estimated ΔP values for Cue A conditional on the presence and absence of B respectively computed after 64 trials. Table 2.4 also presents the estimated ΔP values for Cue A conditional on the presence and absence of B, and Table 2.5 also presents the estimated ΔP values for Cue B conditional on the presence and absence of A. The data are presented for each of the four contingency pairs after 32, 48, and 64 trials. The estimated ΔP data reported in Tables 2.4 and 2.5 correspond to the cumulative values recorded after a given number of trials in that the 32 trial values are based on the first 32 trials, the 48 trial values are based on the first 48 trials, and the 64 trial values are based on

[2]With only four levels of ΔP_B in Experiment 4, the following coefficients were derived to test for a linear trend tracking the ordinal pattern of the conditional ΔP values for A: 22, 18, -12, -28.

Table 2.5: Experiment 4 overall ratings and estimated conditional ΔP values for Cue B after 32, 48, and 64 trials.

		2C-1E							
		0.5/0		0.5/0.25		0.5/0.75		0.5/1	
32	Rating	-22.1	(7.5)	-0.7	(8.8)	61.5	(6.3)	84.3	(5.7)
	$est\Delta P_{B\|A}$	0.12	(0.06)	0.08	(0.05)	0.44	(0.08)	0.68	(0.05)
	$est\Delta P_{B\|\sim A}$	0.15	(0.07)	0.33	(0.08)	0.58	(0.09)	0.75	(0.07)
48	Rating	-10.5	(11.2)	9.3	(9.5)	46.1	(10.7)	90.7	(4)
	$est\Delta P_{B\|A}$	0.15	(0.04)	0.11	(0.04)	0.5	(0.06)	0.79	(0.02)
	$est\Delta P_{B\|\sim A}$	0.12	(0.06)	0.3	(0.09)	0.6	(0.07)	0.8	(0.06)
64	Rating	5.6	(10.4)	-14.6	(7.9)	62.1	(5.3)	87.8	(5.2)
	$est\Delta P_{B\|A}$	0.12	(0.03)	0.14	(0.04)	0.56	(0.05)	0.83	(0.03)
	$est\Delta P_{B\|\sim A}$	0.12	(0.05)	0.26	(0.07)	0.61	(0.07)	0.83	(0.05)
		2E-1C							
		0.5/0		0.5/0.25		0.5/0.75		0.5/1	
32	Rating	-5.6	(9.2)	-3.8	(10.2)	49.9	(11.1)	81.1	(8.6)
	$est\Delta P_{B\|A}$	0.03	(0.06)	0.1	(0.06)	0.46	(0.08)	0.61	(0.07)
	$est\Delta P_{B\|\sim A}$	0.12	(0.08)	0.28	(0.07)	0.63	(0.08)	0.8	(0.07)
48	Rating	-2.1	(8.5)	15	(8.9)	62.8	(8)	69.4	(11.9)
	$est\Delta P_{B\|A}$	0.01	(0.05)	0.1	(0.06)	0.54	(0.07)	0.73	(0.04)
	$est\Delta P_{B\|\sim A}$	0.14	(0.07)	0.22	(0.06)	0.7	(0.06)	0.83	(0.07)
64	Rating	-8.8	(9.5)	13.6	(9.4)	58.9	(8.1)	73	(10.9)
	$est\Delta P_{B\|A}$	0.03	(0.04)	0.09	(0.05)	0.54	(0.06)	0.77	(0.04)
	$est\Delta P_{B\|\sim A}$	0.11	(0.05)	0.22	(0.05)	0.7	(0.06)	0.85	(0.06)

Note. Standard errors of the means are given in parentheses.

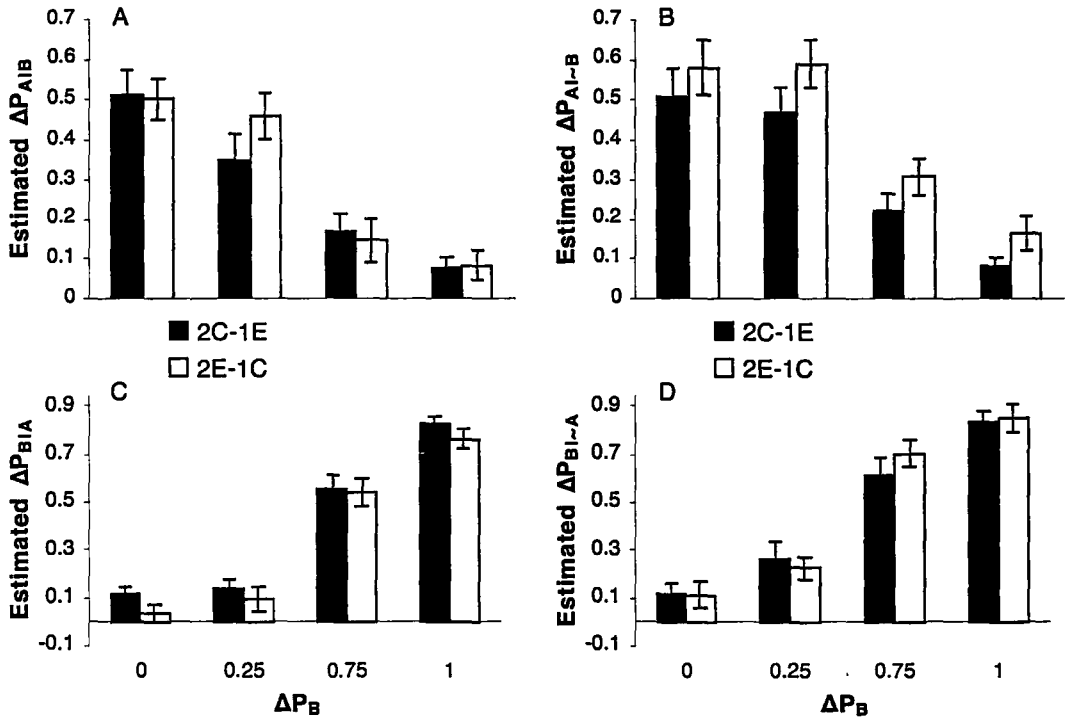

Figure 2.8: Mean estimated conditional ΔP values in Experiment 4 for Events A and B as a function of ΔP_B separately for the two causal scenarios after 64 trials. $est\Delta P_{A|B}$ is shown in Figure 2.8a, $est\Delta P_{A|\sim B}$ is shown in Figure 2.8b, $est\Delta P_{B|A}$ is shown in Figure 2.8c, and $est\Delta P_{B|\sim A}$ is shown in Figure 2.8d. Error-bars represent standard errors of the means.

all of the trials.

The mean estimated conditional ΔP values for A calculated after 32, 48 and 64 trials closely track the conditional ΔP values presented in Table 2.3 for both causal scenarios. Also, the estimated ΔP values conditional on the presence of B ($est\Delta P_{A|B}$) are lower than the estimated ΔP values conditional on the absence of B ($est\Delta P_{A|\sim B}$).

A 2 (scenario: 2C-1E, 2E-1C) × 4 (ΔP_B: 0, 0.25, 0.75, 1) × 3 (trial: 32, 48, 64) × 2 (Cue B status: present, absent) mixed ANOVA on the estimated values for A verifies these observations. The ΔP_B main effect was significant, $F(3, 114) = 33.71$, $MSe = .21$, and contributed to a significant interaction with trial, $F(6, 228) = 9.39$, $MSe = .01$. The status of Cue B main effect was also significant, $F(1, 38) = 19.57$, $MSe = .14$, indicating that estimated conditional ΔP for A was lower when B was present (.28) than when it was absent (.38), and the significant Cue B status × trial interaction, $F(2, 76) = 5.76$, $MSe = .14$, indicates that this difference became less evident across trials. A linear trend analysis was conducted separately for the 2C-1E and 2E-1C causal scenarios after 32, 48, and 64 trials, both on the estimated ΔP values conditional on the presence and absence of B. These analyses reveal a significant linear trend for the prediction responses in both causal scenarios, after each of the three trial intervals (32, 48, 64), regardless of the status of Cue B. Cue-interaction is evident in the prediction responses regardless of the circumstances.

Figures 2.6c and 2.6d illustrate the estimated ΔP values for B conditional on the presence and absence of A respectively computed after 64 trials. An identical ANOVA was performed on the prediction response data for B, substituting Cue A status: (present, absent) for Cue B status. Resembling the data reported in Experiment 3, significant main effects were obtained for ΔP_B, $F(3, 114) = 110.28$, $MSe = .22$, and Cue A status, $F(1, 38) = 19.57$, $MSe = .15$. In addition, the trial factor introduced in Experiment 4 was significant, $F(2, 76) = 9.66$, $MSe = .02$, and led to significant interactions with ΔP_B, $F(6, 228) = 6.93$, $MSe = .01$, and Cue A status, $F(2, 76) = 5.76$, $MSe = .02$. As indicated by the Cue A data, the estimated conditional ΔP values for B were lower when A was present (.36) than when it was absent (.46), and this difference became less evident across trials.

Discussion

The rating data from Experiment 4 are similar to those obtained in each of the previous experiments, and have extended these findings to reveal an interesting scenario × trial interaction. Experiment 4 has shown that cue-interaction is evident across the entire span of 64 trials when A and B are described as two causes of a single effect (2C-1E). When the causal labels are reversed, however, and A and B are described as two effects resulting from a single cause (2E-1C), then we see a very different pattern of results across trials. As in each of the previous experiments, ratings of A reveal that cue-interaction is not evident in the 2E-1C scenario at 32 trials. After 48 and 64 trials, the cue-interaction effect becomes increasingly evident. After 64 trials, ratings of A in the 2E-1C scenario are clearly attenuated, as indicated in Figure 2.7a. While the trial by scenario data in Experiment 1 tended in the same direction as Experiment 4, the effect was not significant. This trial by scenario interaction may not have been evident in Experiment 1 between 32 and 48 trials as we compared the trend between

two trial points (32 and 48) as opposed to three (32, 48, and 64) in Experiment 4. Other data collected in our lab suggest that the trial by scenario interaction is indeed robust (Sadeghi, 2003).

The prediction response values are estimated by separately calculating ΔP conditional on the presence and absence of the other cue. In both Experiment 3 and 4, estimated ΔP conditional on the present cue was significantly lower than estimated ΔP conditional on the absent cue, i.e., $est\Delta P_{A|B} < est\Delta P_{A|\sim B}$ and $est\Delta P_{B|A} < est\Delta P_{B|\sim A}$. Although the conditional ΔP account has not explicitly addressed the relationship between estimated conditional ΔP and actual ΔP, one would expect them to be congruent as indicated by the identical conditional ΔP values presented in Table 2.3. Our data indicating that the estimated values are not congruent with the actual values might be problematic for the conditional ΔP account (see also Tangen & Allan, 2003).

In summary, Experiment 4 provides similar results as Experiments 1-3. Overall ratings were influenced by the causal description of the events after 32 trials. Trial-by-trial prediction responses, however, were not influenced by the causal description of the events. In addition, Experiment 4 demonstrates that on later trials, participants become less sensitive to the difference in description of the two causal scenarios. These data support the argument that causal assessments are not driven solely by associative or causal-model processes, but instead seem be sensitive to both depending on how and when they are obtained. After repeatedly making trial-by-trial predictions, participants may be disregarding the causal order of the events which may be reflected in their overall causal ratings. By continually predicting the presence or absence of the outcome, it is likely that participants are treating the events less like causes and effects, and more like cues and outcomes. As a consequence, on later trials, their causal assessments are based on the same associative strength as their trial-by-trial predictions are based on.

2.8 General Discussion

Price and Yates (1995) were among the first to suggest that both high and low-level processes are used in causal assessments (see also Hagmayer & Waldmann, 2000, for a similar two-process position). There has been little work since then to explain the conditions under which these two processes are likely to be operating. Instead, there has been considerable debate between causal-model and associative learning theorists as to which of the two theoretical interpretations is correct. The results from our experiments revisit the arguments made by Price and Yates (1995) as to the joint contribution of associative and causal factors in judgments of causality.

A similar approach has been taken recently by Collins and Shanks (2002) to account for primacy and recency effects. They describe two strategies involved in judgments of causality: the momentary strategy where judgments simply reflect the current associative strength of the cue, and the integrative strategy where participants do not constrain their judgments on the current perception of the relationship, but instead integrate information across a number of trials. Although Collins and Shanks were describing judgment strategies in primacy and

recency effects, we believe the same tactics are being used in judgments of causally asymmetric events. Participants are required to estimate the presence or absence of the outcome in their trial-by-trial prediction responses. They likely manage this task by identifying cues as generic events without any deeper recognition of their causal status, thereby basing their judgment on the current level of associative strength. Overall ratings, on the other hand, require a more global (integrative) strategy where participants not only consider the status of a single outcome, but also take into account the causal structure of the events presented.

We have demonstrated that the contribution of causal and associative processes depends on what the participant is being asked about the events, and on their experience with those events. Participants recognize that in order to assess the influence of a given cause, they must hold constant (conditionalize on) any alternative causes (2C-1E). Conversely, they understand that a single cause can independently influence a number of effects (1C-2E). In associative terms, two cues compete to be associated with a single outcome. Conversely, one cue can be associated with a number of outcomes. These results are not surprising to anyone. In fact, both causal-model theory and the Rescorla-Wagner model make these predictions. The question, then, is whether the events continue to interact or not when the order of the causal labels are reversed (2E-1C and 1E-2C respectively). The Rescorla-Wagner model predicts that the events should be treated identically in either instance, and causal-model theory predicts that the presence of a cue-interaction effect should reverse along with the causal labels.

Experiments 1 and 2 provide evidence that contingency ratings are influenced by the interaction between causal order (CE vs. EC) and the number of cues and outcomes (2-1 vs. 1-2) indicating that participants are sensitive to the structure of the causal relationship. In Experiment 3, we see that predictions, a second measure of causal assessment, are not so easily swayed by the causal structure of the stimuli. Even though participants assess the same causal relationship in either case, they account for the causal description of the events in one instance (i.e., ratings), but not the other (i.e., predictions). Finally, in Experiment 4, we see that the relative weighting of causal and associative factors are not only influenced by the means of assessing causal inference (ratings and predictions), but also by the repeated exposure to the events. We cannot argue whether the repeated exposure to trial-by-trial predictions is influencing their causal judgments, or whether it is simply the result of additional trials, as these two factors were not tested independently. Regardless, most experiments that support an associative account use *both* a large number of trials and trial-by-trial predictions which may explain the discrepant results. The relative contribution of each of these factors remains an open question.

We would expect that if participants were asked to describe how the causal events were interconnected, or were required to use the causal model for some particular purpose, then they would likely be more sensitive to the structure of the causal relationship then if they were asked to report the probability, covariation, or frequency of the events. Similarly, we might expect participants to consider the causal nature of the events more carefully if several types of causal relationships were presented rather than repeatedly presenting just one. As indicated by the results from Experiment 4, participants become less sensitive to the influence of the causal-model in both their ratings and predictions as trials progress. One might expect that participants would disregard the causal order of the events if they were presented with a large

number of trials. In fact, several experiments supporting an associative interpretation have shown just that. For example, Cobos et al. (2002) required participants to provide a single rating of each event after a learning phase that consisted of as many as 240 trials. Our data from Experiment 4 indicate that any causal-model effect would be largely eliminated by then. While there is no reason to expect the effect of the causal-model to diminish over trials, it may be a step forward in understanding the circumstances under which we use them. We suggest that the number of trials presented to the participant is an important factor in determining their sensitivity to the structure of the causal relationship. In fact, many experiments that have provided support for causal-model theory have used a smaller number of trials (e.g., Waldmann, 2000, 2001) compared to those supporting an associative account (e.g., Shanks & López, 1996; Cobos et al., 2002). This finding may help explain much of the contradictory data in the literature.

Over the past decade, associative and causal-model theorists have continued to debate whether or not human inferences are guided by causal interpretation. We have described specific circumstances that allow one to find one pattern of results or the other, and we provide evidence for an account in which the two processes operate in conjunction rather than independently.

Chapter 3

Assessing (in)sensitivity to causal asymmetry: A matter of degree

> And now remains
> That we find out the cause of this effect,
> Or rather say, the cause of this defect,
> For this effect defective comes by cause.
>
> *Shakespeare (Hamlet II.ii.100-4)*

3.1 Preface

On April 25, 2003, I gave a conference presentation at the 2003 Associative Learning Symposium in Gregynog, Wales, UK. Following the conference, I was notified that Erlbaum publishers were putting together an edited volume based on the Gregynog conference, and the 2003 meeting of the Experimental Psychology Society at Exeter. We were invited to submit a chapter for the volume. The third chapter of this dissertation is reproduced from Tangen et al. (invited) and will be submitted as part of the edited volume. This chapter was written as an extension of the work presented in Tangen and Allan (submitted) (Chapter 2).

3.2 Abstract

In assessing the predictiveness of two causes that result in a common effect (2C-1E), human participants judge one cause in light of the other. However, when asked to rate the diagnosticity of two effects that result from a common cause (2E-1C), they have been shown to conditionalise between the effects in some instances, but not in others. While the debate as to whether human inferences are guided by causal description of the events has continued, few researchers have investigated the specific circumstances under which participants conditionalise between events or not. We report three experiments designed to investigate potential

factors suggested to explain participants' change in sensitivity to causally asymmetric events. Experiment 1 examined whether consistently reminding participants of the direction of the causal relationship would influence their causal assessments. We designed Experiment 2 to investigate whether the wording of the test question used to request participants' ratings is an important variable. Finally, Experiment 3 examined whether asking participants to provide an integrative rating would influence their sensitivity to causal directionality. The results suggest that assessments of causally asymmetric events are best explained by the joint contribution of high and low-level processing whereby causal judgements are influenced both by the causal description and by the associative nature of the events.

3.3 Introduction

As humans became the subject of associationist paradigms, Waldmann and Holyoak (1992) noted the shortcoming of associative models to encode the asymmetry of causal relationships. Causes influence effects, but effects do not influence causes. According to Waldmann (Waldmann & Holyoak, 1990, 1992, 1997; Waldmann, 2000, 2001), associative models neglect the causal status among events by simply encoding the antecedent events as cues and subsequent events as outcomes. For example, in Figure 3.1, the generic events A1, A2, and A3 can be interpreted either as causes or effects. Waldmann argues that humans will rate the influence of each cue differently depending on their causal interpretation and how the cues are interconnected. Specifically, if A1 and A2 are interpreted as two causes that jointly influence a common effect (A3), then according to a common-effect model (2C-1E), one should consider each cause conditional upon the other resulting in an interaction between A1 and A2. On the other hand, if the causal arrows are reversed, and A3 is interpreted as a common cause of two effects (A1 and A2), then according to a common-cause model (2E-1C), the two effects should be considered independently while assessing the causal strength of A3. Both effects are the *product* of the common cause. Therefore, when one is asked to rate the effectiveness of the cause, because A1 and A2 are effects, they have no causal influence on A3. One should, therefore, estimate the unconditional influence on each effect resulting in no cue-interaction. Waldmann maintains that associationist learning theories[1] predict cue-interaction between cues regardless of their causal status, while causal-model theory predicts cue-interaction only between causes.

The data reported by Waldmann and Holyoak (1992) prompted a response among several associative learning theorists who questioned the results and disputed the necessity of causal-models in explaining the presence and absence of cue-interaction. The contentious result was the absence of cue-interaction when two effects preceded a single cause (2E-1C). Because identical cues were used in both the 2C-1E and 2E-1C scenarios, associative models predict an attenuated response in both conditions. In support of associative models, Shanks and López (1996), Matute et al. (1996), and Price and Yates (1995) using different cue-interaction paradigms with various materials, reported a cue-interaction effect regardless of whether two causes preceded a single effect or whether two effects preceded a single cause. Van Hamme

[1]This chapter will focus on the Rescorla-Wagner model as an example of associative models.

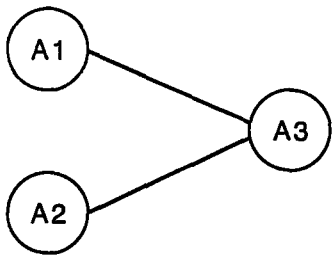

Figure 3.1: Generic casual structure among three interconnected events. A1, A2, and A3 each represent a cause or an effect.

et al. (1993), however, obtained results that were consistent with causal-model theory in which causes interact and effects do not. The authors, however, did not consider their results to be inconsistent with associative theories (see Waldmann, 2000, for review). Waldmann (2000) responded to the criticisms made by the associative theorists by replicating no cue-interaction in the 2E-1C scenario using a novel design to address the criticisms raised against the experiments in Waldmann and Holyoak's (1992) article. He also demonstrated the same sensitivity to causal asymmetry using a one-phase overshadowing design (Waldmann, 2001), Simpson's paradox (Waldmann & Hagmayer, 2001), and he generalised causal-model theory to human categorisation (Waldmann, Holyoak, & Fratianne, 1995; Waldmann & Hagmayer, 1999). In response, Cobos et al. (2002) improved upon the methodology used by Shanks and López (1996) and presented a series of experiments in an attempt to address each of the criticisms proposed by Waldmann (2000, 2001). The results from their analyses reaffirm their previous findings (Shanks & López, 1996) of cue-interaction in the 2E-1C scenario in which multiple effects indicate the presence of a common cause. Cobos et al. (2002) argue that causal asymmetry does not influence the acquisition and use of inferential knowledge.

Recently, however, Tangen and Allan (submitted) provided evidence for both high-level (causal reasoning) processes, and low-level (associative) processes. They argued that both factors influence causal assessment depending on what is being asked about the events, and participants' experience with those events. In particular, in two experiments, they demonstrated how expectations of the structure of causal relationships influence overall causal ratings. In two other experiments, they showed that participants are insensitive to causal structure in their trial-by-trial prediction responses in contrast to their overall ratings, and that the effect of expectation on overall ratings are attenuated by including prediction responses and increasing the total number of trials presented. Tangen and Allan (submitted) concluded that people engage in both causal reasoning and associative learning.

Several potential factors have been suggested to explain the presence or absence of a causal model effect. The present chapter will examine the influence of three such factors in the one-phase simultaneous blocking design used by Tangen and Allan (submitted). The following experiments were not designed to set associative and causal-model theories in opposition (see Tangen & Allan, submitted, for contrast), but rather to examine three circumstances under which a causal model effect occurs or not, and to further examine the role of conditionalisation

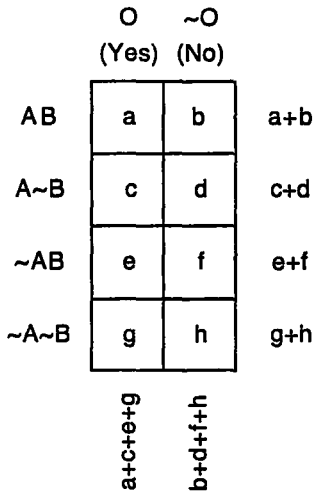

Figure 3.2: 4×2 contingency matrix illustrating the eight possible cue-outcome combinations for two cues. Each cell represents the frequency of each event type.

in causal-model theory.

3.4 Causal-Model Theory and Conditional ΔP

The one-phase simultaneous blocking design initially proposed by Baker et al. (1993) makes use of all possible event combinations for two cues (A and B) and a common outcome (O): both cues may be present (AB), one may be present and the other absent (A~B or ~AB), or both may be absent (~A~B) as illustrated in Figure 3.2. For each cue combination, the outcome either occurs (O) or not (~O). In such a task, after being presented with a cue combination, participants are typically asked to predict whether the outcome will occur or not (Yes, No), which is also represented in Figure 3.2. By calculating the frequency that each event combination occurred, Cues A and B can be expressed in terms of their unconditional

or conditional ΔP values respectively:

$$\Delta P_A = P(O|A) - P(O|\sim A) = \frac{a+c}{a+b+c+d} - \frac{e+g}{e+f+g+h} \tag{3.1}$$

$$\Delta P_B = P(O|B) - P(O|\sim B) = \frac{a+e}{a+b+e+f} - \frac{c+g}{c+d+g+h} \tag{3.2}$$

$$\Delta P_{A|B} = P(O|AB) - P(O|\sim AB) = \frac{a}{a+b} - \frac{e}{e+f} \tag{3.3}$$

$$\Delta P_{A|\sim B} = P(O|A\sim B) - P(O|\sim A\sim B) = \frac{c}{c+d} - \frac{g}{g+h} \tag{3.4}$$

$$\Delta P_{B|A} = P(O|BA) - P(O|\sim BA) = \frac{a}{a+b} - \frac{c}{c+d} \tag{3.5}$$

$$\Delta P_{B|\sim A} = P(O|B\sim A) - P(O|\sim B\sim A) = \frac{e}{e+f} - \frac{g}{g+h} \tag{3.6}$$

in which the two unconditional ΔP values (Equations 3.1 and 3.2) correspond to the difference between the proportion of times the outcome occurs given the cue and the proportion of times the outcome occurs not given the cue. The conditional ΔP values in Equations 3.3 - 3.6 allow one to assess the influence of each cue both in the presence and absence of the other cue.

Cue-interaction by means of conditionalisation occurs if participants' ratings better correspond with conditional, rather than unconditional, ΔP. When two cues described as causes precede a single outcome described as an effect (2C-1E), a causal-model account predicts that participants should conditionalise by rating the influence of each cause relative to the other. Associative models also predict that the cues should interact but by means of cue competition as opposed to conditionalisation. Alternatively, if the two cues are described as effects while the outcome is described as a cause (2E-1C), a causal-model account predicts that participants should rate each effect independently coinciding with unconditional ΔP described above. Because associative accounts disregard the causal description of the events, they would again predict that the cues would interact by means of cue competition.

Each of the experiments described in Tangen and Allan (submitted) were designed to set apart the two accounts of cue-interaction. In each case, two cues were paired with a single outcome where Cue A was always moderately contingent ($\Delta P_A = 0.5$), and was paired with Cue B which ranged from being non-contingent ($\Delta P_B = 0$) to perfectly contingent ($\Delta P_B = 1$). If participants' ratings or predictions of A changed as the contingency of B increased, then they were judging A relative to B. In contrast, if their judgement of A did not change as B increased, then they were judging A independently of B. Throughout each of the four experiments, Tangen and Allan found that on early trials there was a significant causal model effect in participants' ratings where the causal scenario (2C-1E and 2E-1C) significantly interacted with Cue B contingency. In Experiment 4, they showed that the causal model effect dissipated as the trials progressed. Furthermore, they revealed a dissociation between ratings

and prediction responses where participants' ratings were sensitive to the causal scenario (at least on early trials), while their prediction responses were insensitive to the causal scenario instead reflecting what seems to be the current level of associative strength.

3.5 Alleged Causal Model Influences

The debate between associative and causal-model theorists prompted a number of experiments examining the influence of causal asymmetry on judgements of contingency, and has resulted in a number of positive and negative results (Waldmann & Holyoak, 1992; Van Hamme et al., 1993; Price & Yates, 1995; Shanks & López, 1996; Waldmann, 2001; Cobos et al., 2002; Tangen & Allan, submitted). Several potential factors have been suggested to explain the presence or absence of a causal model effect. We will examine the influence of three such factors using the one-phase blocking design described above. Recall that in Tangen and Allan (submitted, Experiment 4), participants became less sensitive to the causal direction of the events as trials progressed. The purpose of the present series of experiments is to examine the specific circumstances under which we can alter participants' sensitivity to the direction of the causal relationship. In order to determine whether each of the alleged causal model influences was effective in bringing about a causal model effect, we will compare each experiment in turn with the results obtained in Experiment 4 of Tangen and Allan (submitted). Therefore, using those data as a template, we will compare them with the data obtained from each experiment in sequence by including both in an analysis of variance (ANOVA), and we will report any significant differences between them. However, we are interested primarily in whether there is any change in participants' sensitivity to the causal model for each of the following experiments as measured by a significant experiment × scenario × ΔP_B interaction.

3.5.1 General Method

Participants and Design

A total of 160 undergraduate students at McMaster University volunteered for course credit (40 participants in Experiments 1 and 3, and 80 participants in Experiment 2). All of the experiments used the same one-phase blocking design described in Tangen and Allan (submitted, Experiment 4). Sixty-four trials were presented to each participant. On each trial, one of four possible cue-combinations was presented (AB, A∼B, ∼AB, ∼A∼B) at which point participants were asked to predict whether the outcome on that particular trial would occur or not (Yes, No). The actual outcome of the trial (O, ∼O) followed along with corrective feedback on their decision (Correct, Incorrect). Eight trial types were therefore possible and are presented in Table 3.1. The number of times that each trial type was presented was determined by pairing moderately contingent Cue A ($\Delta P_A = 0.5$) with Cue B which varied in contingency (ΔP_B: 0, 0.25, 0.75, 1). The particular event frequencies were selected so that the conditional ΔP_A values would gradually diverge from the unconditional ΔP_A values as ΔP_B increased.

Cue-interaction was measured by the pattern of data obtained from ratings of Cue A:

Table 3.1: Frequency of events in Experiment 1-3. Unconditional ΔP values were calculated using Equations 3.1 and 3.2. Conditional ΔP values were calculated using Equations 3.3-3.6.

Trial Type	0.5/0	0.5/0.25	0.5/0.75	0.5/1
ABO	12	16	22	24
A∼BO	12	8	2	0
∼ABO	4	4	6	8
∼A∼BO	4	4	2	0
AB∼O	4	4	2	0
A∼B∼O	4	4	6	8
∼AB∼O	12	8	2	0
∼A∼B∼O	12	16	22	24
# of Trials	64	64	64	64
ΔP_A	0.5	0.5	0.5	0.5
$\Delta P_{A\|B}$	0.5	0.47	0.17	0
$\Delta P_{A\|\sim B}$	0.5	0.47	0.17	0
ΔP_B	0	0.25	0.75	1
$\Delta P_{B\|A}$	0	0.13	0.67	1
$\Delta P_{B\|\sim A}$	0	0.13	0.67	1

if ratings of A changed as a function of ΔP_B, then the cues interacted, if ratings of A did not vary, then the cues did not interact. According to causal-model theory, by means of conditional ΔP, participants in the 2C-1E scenario should rate A conditional on B thereby tracking the pattern of conditional ΔP_A values across the four contingency pairs (0.5, 0.47, 0.17, 0). Participants in the 2E-1C scenario should rate A independently of B thereby tracking the pattern of unconditional ΔP_A values across the four contingency pairs (0.5, 0.5, 0.5, 0.5). Half of the participants were therefore assigned to the 2C-1E scenario and half were assigned to the 2E-1C scenario, and the four contingency pairs were presented to each participant in random order.

Throughout each set of 64 trials, participants were asked to rate Cues A and B after 16, 32, 48, and 64 trials. When participants were asked to predict the outcome of each trial, the cue combinations and their corresponding predictions (Yes, No) were mapped onto the 4×2 matrix presented in Figure 3.2 and used as an indirect measure of participants' conditional ΔP estimates (López et al., 1998; Tangen & Allan, 2003).

Procedure and Materials

Participants received instructions on a computer screen where they were informed about four strains of bacteria that have been discovered in the mammalian digestive system. In the 2C-1E scenario, they were told that scientists were testing whether a pair of chemicals affected the strain's survival, whereas, in the 2E-1C scenario, the scientists were testing whether the bacteria affected the production of a pair of chemicals.

Up to four participants at a time performed the experiment on Power Macintosh computers. Each experiment was programmed in MetaCard 2.4.3. In the instructions, the four contingency pairs were identified as separate "experiments" to test the influence of the chemicals on the bacteria, or vice versa. Within each causal scenario, the 64 trials were presented in random order according to the frequencies presented in Table 3.1. The addition or production of a chemical was indicated by a computer rendered movie of a coloured three-dimensional chemical spinning along its axis, and actual footage of moving bacteria was displayed when the bacteria survived or were added. Faded, unmoving greyscale images of the same chemicals and bacteria were displayed to indicate their absence on a given trial. The names of the chemicals and bacteria were displayed only when the events occurred. Each of the movies and images were randomly assigned fictitious names from a set of eight chemicals and four bacteria, and were randomly assigned to appear on the left or right-hand side of the screen for each of the 64 trials.

Participants were presented with one of four cue combinations consisting of the presence or absence of two chemicals and were then asked to indicate whether they thought the bacterial strain survived/was added or not by clicking one of two buttons on the computer screen. Once they made their selection, they were presented with the outcome along with *Correct* or *Incorrect* as feedback. After viewing and predicting the outcome for 16 trials, participants in the 2C-1E causal scenario were asked to rate how strongly each chemical (Cues A and B) affected the bacteria, and those in the 2E-1C scenario were asked to rate how strongly the bacteria affected the production of each chemical (Cues A and B). Ratings were made on a

scale ranging from -100 to 100 by moving a horizontal scrollbar with a mouse ranging from -100 at the leftmost position to 100 at the rightmost position, anchored at 0 at the centre. The trials resumed and participants were asked to repeat the rating process after 32, 48, and 64 trials.

3.6 Experiment 1: Clarity of the Causal Model

Waldmann and Holyoak (1997) presented a list of methodological requirements for investigating causal directionality, the first of which is to "ensure that participants consistently interpret the learning situation in terms of directed cause-effect relations" (p. 127). In criticising the methodology used by Shanks and López (1996), Waldmann and Holyoak insisted on the necessity of making the causal relationship in the instructions and materials unmistakable. They argued that if there is any ambiguity in participants' interpretation of the event relations, then one cannot accurately measure the influence of their causal interpretation on learning. While Shanks and López used a cover story in which various symptoms (effects) were diagnostic of a disease (cause), Waldmann and Holyoak argued that the directionality of the causal relationship could be misinterpreted.

Tangen and Allan (submitted) used materials that could also be interpreted as causally ambiguous. Participants were told that scientists have recently discovered several strains of bacteria that exist in the mammalian digestive system. In the 2C-1E scenario, they were told that the scientists were testing whether certain pairs of chemicals (causes) affect the survival of the bacteria (effects). In the 2E-1C scenario, they were told that the scientists were testing whether certain pairs of chemicals (effects) were produced as a result of the addition of the bacteria (cause). Tangen and Allan selected materials particularly for the potential to reverse causal direction, in which the chemicals could be described as causes just as easily as effects. In their final experiment, Tangen and Allan demonstrated that the influence of the causal model decreased as the number of trials increased. It is possible that over time, the causal relationship between the events became less clear where participants misinterpreted the cues as causes and the outcome as an effect in the 2E-1C scenario. In their methodological requirements, Waldmann and Holyoak (1997) made a point of noting that participants must *consistently* interpret the learning situation in terms of directed cause-effect relations. Experiment 1 was designed to serve this purpose by consistently reminding participants about the direction of the causal relationship at hand.

3.6.1 Method

Participants in Experiment 1 were asked to provide contingency ratings for Cues A and B after 16, 32, 48, and 64 trials. Between ratings, they were exposed to the sixteen causal model prompts given in the Appendix to consistently remind them about the direction and nature of the causal relationship. Those in the 2C-1E condition were given 2C-1E prompts while those in the 2E-1C condition were given 2E-1C prompts. Each participant was therefore presented with each of the 16 prompts two times, in random order, after Trials 8, 16, 24, 32, 40, 48, 56, and 64 for each of the four contingency pairs shown in Table 3.1.

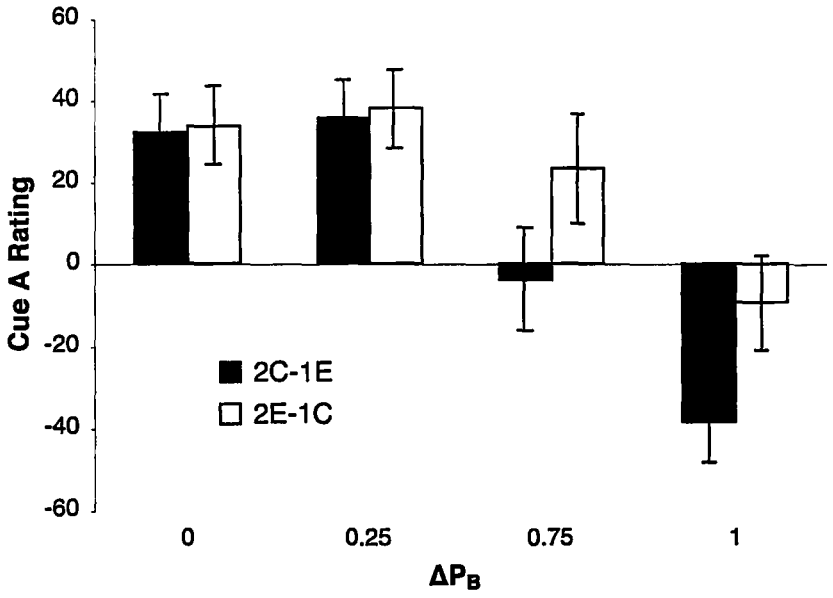

Figure 3.3: Mean ratings in Experiment 1 after 64 trials of Cue A (Figure 3.3). The ratings are shown as a function of ΔP_B (0, 0.25, 0.75, 1) separately for each of the two conditions. Error-bars represent standard errors of the means.

3.6.2 Results[2]

Ratings

Figure 3.3 illustrates the mean ratings of Cue A after 64 trials, and Table 3.2 provides the mean ratings of Cue A after 32, 48, and 64 trials.[3] The data are presented for each of the four contingency pairs. The pattern of data for ratings of Cue A is virtually identical to the results obtained in Tangen and Allan (submitted, Experiment 4). Ratings of Cue A decrease as ΔP_B increases closely tracking the pattern of conditional ΔP values presented in Table 3.1 regardless of whether the two cues were described as causes (2C-1E) or effects (2E-1C).

The data from Tangen and Allan (submitted, Experiment 4) was used as a template

[2]Several different frequencies can be selected to fill the eight cells of the 4×2 matrix each resulting in various combinations of unconditional and conditional ΔP values. The frequencies shown in Table 1 were selected to produce a descending pattern of conditional ΔP_A values while maintaining identical unconditional ΔP_A values. As well, they were selected so the unconditional and conditional ΔP_B values would be as closely matched as possible. The ΔP_B values were therefore selected only for their influence on the conditional ΔP_A values. The data for Cue B are virtually identical in each of the three experiments reported here, and do not differ from Experiments 1-4 in Tangen & Allan (under review). Therefore, the data for Cue B will not be reported in this chapter.

[3]The ratings and estimated ΔP values after 16 trials are not reported as participants' prediction responses of the randomly presented events occasionally resulted in 4×2 matrices with row frequencies of zero.

Table 3.2: Experiment 1 mean ratings and estimated ΔP values of Cue A conditional on the presence of Cue B ($est\Delta P_{A|B}$) and the absence of Cue B ($est\Delta P_{A|\sim B}$) after 32, 48, and 64 trials.

		2C-1E							
		0.5/0		0.5/0.25		0.5/0.75		0.5/1	
32	Rating	39.7	(8.6)	29.5	(11.5)	-16	(12.6)	-39.7	(9.7)
	$est\Delta P_{A\|B}$	0.39	(0.08)	0.3	(0.09)	0.19	(0.07)	0.09	(0.06)
	$est\Delta P_{A\|\sim B}$	0.46	(0.08)	0.38	(0.09)	0.37	(0.05)	0.16	(0.04)
48	Rating	25.3	(9.8)	26.6	(9.8)	0.2	(12)	-45.7	(10.4)
	$est\Delta P_{A\|B}$	0.38	(0.07)	0.36	(0.08)	0.16	(0.05)	0.07	(0.03)
	$est\Delta P_{A\|\sim B}$	0.5	(0.07)	0.44	(0.07)	0.29	(0.05)	0.09	(0.03)
64	Rating	32.5	(9)	35.6	(9.4)	-3.7	(12.4)	-38.5	(9.6)
	$est\Delta P_{A\|B}$	0.39	(0.07)	0.42	(0.07)	0.13	(0.04)	0.06	(0.02)
	$est\Delta P_{A\|\sim B}$	0.51	(0.07)	0.48	(0.06)	0.27	(0.05)	0.08	(0.02)
		2E-1C							
		0.5/0		0.5/0.25		0.5/0.75		0.5/1	
32	Rating	35.4	(11.1)	45.6	(10.1)	11.4	(10.8)	0.1	(13.9)
	$est\Delta P_{A\|B}$	0.39	(0.08)	0.59	(0.05)	0.3	(0.07)	0.15	(0.06)
	$est\Delta P_{A\|\sim B}$	0.61	(0.05)	0.61	(0.08)	0.44	(0.08)	0.23	(0.07)
48	Rating	27.4	(11.9)	50.8	(5.8)	28.1	(10.3)	-1.6	(14.1)
	$est\Delta P_{A\|B}$	0.46	(0.07)	0.54	(0.05)	0.32	(0.06)	0.09	(0.04)
	$est\Delta P_{A\|\sim B}$	0.62	(0.06)	0.6	(0.08)	0.38	(0.07)	0.16	(0.06)
64	Rating	33.9	(9.5)	38	(9.5)	23.2	(13.3)	-9.4	(11.6)
	$est\Delta P_{A\|B}$	0.48	(0.07)	0.54	(0.06)	0.27	(0.05)	0.07	(0.03)
	$est\Delta P_{A\|\sim B}$	0.64	(0.05)	0.59	(0.07)	0.37	(0.06)	0.11	(0.04)

Note. Standard errors of the means are given in parentheses.

to compare the results from the present experiment and two subsequent experiments. In Tangen and Allan (submitted, Experiment 4), at 32 trials, the pattern of results was similar to that of the previous three experiments. In the 2C-1E scenario, ratings of A tracked the pattern of conditional ΔP values presented in Table 3.1, whereas in the 2E-1C scenario the ratings tracked the pattern of unconditional ΔP values. After 48 and 64 trials, however, a different pattern of results emerged. Ratings of A declined as ΔP_B increased, regardless of the causal scenario. The effect of the causal-model seemed to have dissipated over trials, and cue-interaction occurred for both scenarios. A 2 (scenario: 2C-1E, 2E-1C) × 4 (ΔP_B: 0, 0.25, 0.75, 1) × 3 (trial: 32, 48, 64) mixed ANOVA on the ratings of A revealed only a significant main effect for ΔP_B, $F(3, 114) = 21.27$, $p < .001$, which contributed to significant interactions with scenario, $F(3, 114) = 3.33$, $p < .05$, trial, $F(6, 228) = 2.60$, $p < .05$, and a three-way interaction with scenario and trial, $F(6, 228) = 3.13$, $p < .01$. The two-way interaction of ΔP_B with scenario indicates that cue-interaction was greater for 2C-1E than for 2E-1C, and the three-way interaction with trial indicates that this difference decreased over trials.

Using these results as a template, a 2 (experiment: Exp 4, Exp 1) × 2 (scenario: 2C-1E, 2E-1C) × 4 (ΔP_B: 0, 0.25, 0.75, 1) × 3 (trial: 32, 48, 64) mixed ANOVA was conducted on the ratings of Cue A. Participants' sensitivity to causal direction did not differ between the two experiments as indicated by the experiment × scenario × ΔP_B interaction which was not significant, $F(3, 228) = .62$, $p > .05$, and the experiment × scenario × ΔP_B × trial interaction which also was not significant, $F(6, 456) = 1.34$, $p > .05$. The only effect (that interacted with experiment) to reach significance was a three-way interaction between experiment, trial, and ΔP_B, $F(6, 456) = 3.05$, $p < .01$.

Predictions

Table 3.2 also provides the mean estimated ΔP values conditional on the presence and absence of Cue B after 32, 48, and 64 trials. The prediction response data is similar to the rating data in that the estimated ΔP_A values decrease as ΔP_B increases regardless of whether the two cues were described as causes (2C-1E) or effects (2E-1C).

In Tangen and Allan (submitted, Experiment 4), the mean estimated conditional ΔP values for A calculated after 32, 48 and 64 trials closely tracked the conditional ΔP values presented in Table 3.1 for both causal scenarios. Also, the estimated ΔP values conditional on the presence of B ($est\Delta P_{A|B}$) were lower than the estimated ΔP values conditional on the absence of B ($est\Delta P_{A|\sim B}$). A 2 (scenario: 2C-1E, 2E-1C) × 4 (ΔP_B: 0, 0.25, 0.75, 1) × 3 (trial: 32, 48, 64) × 2 (Cue B status: present, absent) mixed ANOVA on the estimated values for A revealed a significant main effect of ΔP_B, $F(3, 114) = 33.71$, $p < .001$, and contributed to a significant interaction with trial, $F(6, 228) = 9.39$, $p < .001$. The status of Cue B main effect was also significant, $F(1, 38) = 19.57$, $p < .001$, and interacted with trial, $F(2, 76) = 5.76$, $p < .01$.

Again, using these data as a template, a 2 (experiment: Exp 4, Exp 1) × 2 (scenario: 2C-1E, 2E-1C) × 4 (ΔP_B: 0, 0.25, 0.75, 1) × 3 (trial: 32, 48, 64) × 2 (Cue B status: present, absent) mixed ANOVA was conducted on the estimated conditional ΔP values for Cue A. As in the rating data, participants' sensitivity to causal direction did not differ between the

two experiments as indicated by the experiment × scenario × ΔP_B interaction which was not significant, $F(3, 228) = .15$, $p > .05$. The only effect (that interacted with experiment) to reach significance in the prediction data was a four-way interaction between experiment, scenario, ΔP_B, and trial, $F(6, 456) = 2.97$, $p < .01$.

3.6.3 Discussion

Experiment 1 tested the methodological requirement proposed by Waldmann and Holyoak (1997) for clarity of the causal model by consistently reminding participants about the direction and nature of the causal relationship for two causal scenarios. The results indicate that the frequent presentation of causal prompts did not influence their sensitivity to causal directionality.[4] Neither participants' ratings or predictions differed between the data obtained from the present experiment and Tangen and Allan (submitted, Experiment 4). Therefore, the first alleged causal model influence of emphasising the clarity of the causal model was ineffective.

3.7 Experiment 2: Test Question

Cue-interaction, according to Matute et al. (1996), depends on the nature of the test question. Specifically, when presented with two cues that share a common outcome (as illustrated in Figure 3.1), Matute et al. argue that asking about the *contiguity* of events will result in no cue-interaction regardless of whether the events are described as causes or effects. The contiguous test question is thought to discourage comparisons between cues resulting in judgements that are based on unconditional ΔP (Equations 3.1 and 3.2) rather than conditional ΔP (Equations 3.3 - 3.6). In contrast, asking about the *causality* of events will result in cue-interaction. The causal test question is thought to encourage comparisons between causes when potential alternative causes are present, but not when only one potential cause is present. Therefore, judgements should be based on conditional ΔP rather than unconditional ΔP in the 2C-1E scenario, but not the 2E-1C scenario.[5] Experiment 2 was designed to examine the predictions made by Matute et al. (1996) by crossing the two causal scenarios tested in Experiment 1 (2C-1E and 2E-1C) with causal and contiguous test questions.

3.7.1 Method

The design, procedure, and materials for Experiment 2 was identical to that of Experiment 1 without the presentation of the causal model prompts introduced in the first experiment.

[4]Reminding participants about the direction and nature of the causal relationship after every 8 trials was actually our second attempt to produce a causal model effect. The first attempt presented the same causal-model prompts after every 16 trials. The results from the two experiments were virtually identical.

[5]Matute et al. (1996) also tested the "directionality" of the test question, i.e., cause-to-effect (CE) or effect-to-cause (EC) questions, though ratings did not seem to be influenced by whether the questions were worded one way or the other when asked questions about causality. Therefore, we will limit our discussion to causal vs. contiguous test questions.

After observing a series of 16 trials, participants in both the 2C-1E and 2E-1C groups were asked either a causal or contiguous test question. Both "test question" and "causal scenario" were between factors resulting in four groups of subjects who were asked one of the following questions:

2C-1E Causal: To what degree do you think placing the chemical in the petri dish caused the bacterial strain to survive?

2C-1E Contiguous: To what degree do you think placing the chemical in the petri dish was related, even by mere chance, to the survival of the bacterial strain?

2E-1C Causal: To what degree do you think placing the bacterial strain in the digestive system caused each chemical to be produced?

2E-1C Contiguous: To what degree do you think placing the bacterial strain in the digestive system was related, even by mere chance, to the production of each chemical?

Matute et al. (1996) asked participants whether "C is the cause of E?" in the causal test question or "when C is present, does E co-occur?" in the contiguous test question. It is not obvious how these questions might map on to a rating scale that ranges from -100 to 100. Therefore, the wording of the causal and contiguous test questions was adapted from Matute, Vegas, and De Marez (2002) to better correspond to the rating scale used in the present series of experiments.

3.7.2 Results

As in Experiment 1, there were two dependent measures, ratings and predictions.

Ratings

Figures 3.4a and 3.4b illustrate the mean ratings of Cue A after 64 trials for the causal and contiguous conditions respectively. Tables 3.3 and 3.4 depict the mean ratings of Cue A for the causal and contiguous groups respectively. The data are presented for each of the four contingency pairs after 32, 48, and 64 trials. The pattern of data is very similar to the results obtained in both Experiment 1 and Tangen and Allan (submitted, Experiment 4) for *both* the causal and contiguous groups. Ratings of Cue A decrease as ΔP_B increases closely tracking the pattern of conditional ΔP values presented in Table 3.1 regardless of whether the two cues were described as causes (2C-1E) or effects (2E-1C) or whether the participants were asked a causal or contiguous test question.

Again, using the data from Tangen and Allan (submitted, Experiment 4) as a template, a 2 (experiment: Exp 4, Exp 2) × 2 (scenario: 2C-1E, 2E-1C) × 4 (ΔP_B: 0, 0.25, 0.75, 1) × 3 (trial: 32, 48, 64) mixed ANOVA was conducted on the ratings of Cue A separately for the two test question conditions (causal, contiguous). As in the previous experiment, there was no effect of the causal model as indicated by the experiment × scenario × ΔP_B interaction which was not significant for the group of participants given the causal test question, $F(3, 228) = .01$, $p > .05$, or those given the contiguous test question, $F(3, 228) = .18$, $p > .05$. The only effect to reach significance that interacted with experiment, was a three-way interaction between experiment, trial, and ΔP_B, $F(6, 456) = 3.05$, $p < .01$, for the causal test question group, in

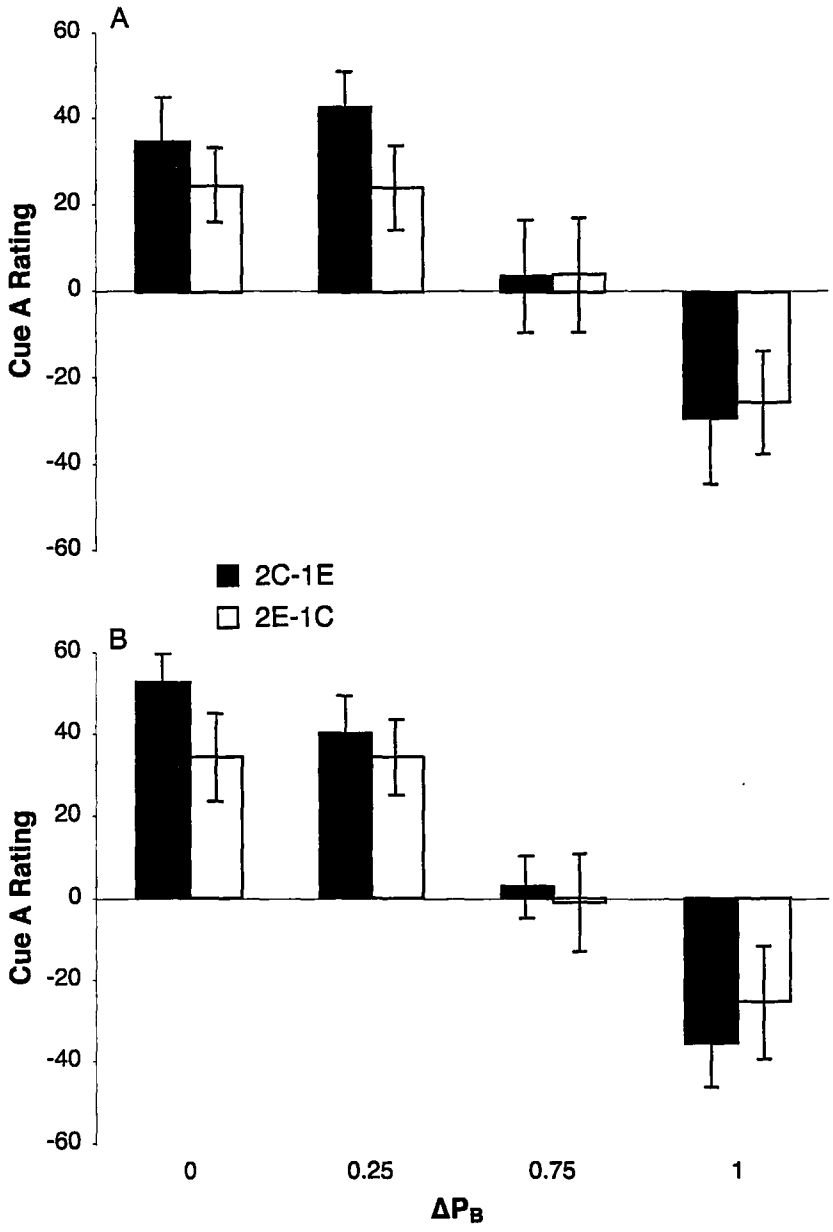

Figure 3.4: Mean ratings in Experiment 2 after 64 trials of Cue A for participants given the *Causal* test question (Figure 3.4a) and the *Contiguous* test question (Figure 3.4b). The ratings are shown as a function of ΔP_B (0, 0.25, 0.75, 1) separately for each of the two conditions. Error-bars represent standard errors of the means.

Table 3.3: Experiment 2 mean ratings and estimated ΔP values of Cue A conditional on the presence of Cue B ($est\Delta P_{A|B}$) and the absence of Cue B ($est\Delta P_{A|\sim B}$) after 32, 48, and 64 trials for the Causal test question.

		2C-1E							
		0.5/0		0.5/0.25		0.5/0.75		0.5/1	
32	Rating	45.4	(9.9)	33.6	(11)	8.7	(14.8)	-41.4	(13.7)
	$est\Delta P_{A\|B}$	0.39	(0.09)	0.25	(0.08)	0.11	(0.07)	0.04	(0.04)
	$est\Delta P_{A\|\sim B}$	0.52	(0.09)	0.45	(0.09)	0.39	(0.08)	0.21	(0.05)
48	Rating	44.1	(10.3)	32.5	(9.3)	26.7	(11.4)	-33.9	(12.5)
	$est\Delta P_{A\|B}$	0.44	(0.07)	0.33	(0.07)	0.15	(0.06)	0.03	(0.04)
	$est\Delta P_{A\|\sim B}$	0.52	(0.08)	0.49	(0.07)	0.4	(0.08)	0.15	(0.05)
64	Rating	34.8	(10.3)	42.7	(8.1)	3.7	(13)	-29.3	(15.2)
	$est\Delta P_{A\|B}$	0.46	(0.08)	0.42	(0.06)	0.14	(0.05)	0.01	(0.03)
	$est\Delta P_{A\|\sim B}$	0.53	(0.08)	0.56	(0.07)	0.38	(0.07)	0.1	(0.04)
		2E-1C							
		0.5/0		0.5/0.25		0.5/0.75		0.5/1	
32	Rating	38.2	(6.5)	33.5	(12.1)	10.5	(13.1)	-7.6	(11.5)
	$est\Delta P_{A\|B}$	0.49	(0.08)	0.43	(0.06)	0.19	(0.05)	0.18	(0.07)
	$est\Delta P_{A\|\sim B}$	0.71	(0.05)	0.6	(0.06)	0.51	(0.07)	0.2	(0.04)
48	Rating	22	(8.6)	29.4	(7.9)	2.4	(10.3)	-16.5	(12.2)
	$est\Delta P_{A\|B}$	0.47	(0.07)	0.46	(0.06)	0.17	(0.05)	0.11	(0.04)
	$est\Delta P_{A\|\sim B}$	0.73	(0.04)	0.63	(0.06)	0.39	(0.04)	0.2	(0.03)
64	Rating	24.6	(8.4)	24.2	(9.7)	3.9	(13)	-25.7	(11.9)
	$est\Delta P_{A\|B}$	0.51	(0.06)	0.42	(0.06)	0.16	(0.04)	0.1	(0.04)
	$est\Delta P_{A\|\sim B}$	0.73	(0.04)	0.63	(0.05)	0.34	(0.04)	0.16	(0.03)

Note. Standard errors of the means are given in parentheses.

Table 3.4: Experiment 2 mean ratings and estimated ΔP values of Cue A conditional on the presence of Cue B ($est\Delta P_{A|B}$) and the absence of Cue B ($est\Delta P_{A|\sim B}$) after 32, 48, and 64 trials for the Contiguous test question.

		2C-1E							
		0.5/0		0.5/0.25		0.5/0.75		0.5/1	
32	Rating	30.1	(11.3)	33.8	(10.4)	11.7	(8.2)	-25.3	(13.9)
	$est\Delta P_{A\|B}$	0.45	(0.06)	0.34	(0.08)	0.12	(0.04)	0.02	(0.03)
	$est\Delta P_{A\|\sim B}$	0.55	(0.08)	0.44	(0.08)	0.47	(0.09)	0.24	(0.08)
48	Rating	46	(8.3)	38.6	(9.8)	2.4	(7.7)	-33.7	(11.1)
	$est\Delta P_{A\|B}$	0.54	(0.06)	0.41	(0.08)	0.15	(0.04)	0.03	(0.02)
	$est\Delta P_{A\|\sim B}$	0.66	(0.08)	0.6	(0.06)	0.43	(0.08)	0.14	(0.06)
64	Rating	52.8	(6.7)	40	(9.5)	2.7	(7.5)	-35.6	(10.6)
	$est\Delta P_{A\|B}$	0.63	(0.04)	0.45	(0.07)	0.12	(0.03)	0.02	(0.02)
	$est\Delta P_{A\|\sim B}$	0.7	(0.07)	0.63	(0.06)	0.37	(0.07)	0.11	(0.05)
		2E-1C							
		0.5/0		0.5/0.25		0.5/0.75		0.5/1	
32	Rating	29.3	(9.1)	26.3	(11.2)	6.6	(12.4)	-4.7	(13.3)
	$est\Delta P_{A\|B}$	0.53	(0.06)	0.53	(0.06)	0.28	(0.08)	0.15	(0.05)
	$est\Delta P_{A\|\sim B}$	0.71	(0.07)	0.62	(0.08)	0.51	(0.08)	0.33	(0.07)
48	Rating	33.3	(9.1)	16.8	(11.4)	-4.7	(12)	-18.6	(12.9)
	$est\Delta P_{A\|B}$	0.55	(0.06)	0.57	(0.05)	0.25	(0.06)	0.08	(0.03)
	$est\Delta P_{A\|\sim B}$	0.72	(0.06)	0.63	(0.07)	0.49	(0.08)	0.24	(0.05)
64	Rating	34.4	(10.7)	34.1	(9.4)	-1.1	(11.9)	-25.5	(13.9)
	$est\Delta P_{A\|B}$	0.58	(0.06)	0.59	(0.05)	0.18	(0.05)	0.06	(0.02)
	$est\Delta P_{A\|\sim B}$	0.73	(0.05)	0.66	(0.06)	0.47	(0.08)	0.19	(0.05)

Note. Standard errors of the means are given in parentheses.

which ratings in the present experiment were slightly higher when $\Delta P_B = 0.75$ during later trials than in Experiment 4. There were no significant interactions with experiment for the contiguous group.

Predictions

Tables 3.3 and 3.4 also depict the mean estimated conditional ΔP values of Cue A conditional on the presence and absence of Cue B for the causal and contiguous groups respectively. The data are presented for each of the four contingency pairs after 32, 48, and 64 trials. The prediction responses are similar to the rating data in that Cue A decreases as ΔP_B increases regardless of whether the two cues were described as causes or effects, or whether participants were asked a causal or contiguous test question.

Again, using the data from Tangen and Allan (submitted, Experiment 4) as a template, a 2 (experiment: Exp 4, Exp 1) × 2 (scenario: 2C-1E, 2E-1C) × 4 (ΔP_B: 0, 0.25, 0.75, 1) × 3 (trial: 32, 48, 64) × 2 (Cue B status: present, absent) mixed ANOVA was conducted on the estimated conditional ΔP values for Cue A separately for the two test questions (causal, contiguous). As in the rating data, there was no significant effect of the causal model between the two experiments as indicated by the experiment × scenario × ΔP_B interaction which was not significant either for participants in the causal group, $F(3, 228) = .11$, $p > .05$, or the contiguous group, $F(3, 228) = .13$, $p > .05$. For those given the causal test question, the only effects to interact with experiment was a significant three-way interaction between experiment, scenario, and trial, $F(2, 152) = 3.52$, $p < .05$. For those given the contiguous test question, the only significant effects to interact with experiment were trial, $F(2, 152) = 3.22$, $p < .05$, and trial × scenario, $F(2, 152) = 6.83$, $p < .01$. The experiment main effect was also significant for the contiguous group indicating that the prediction reponses were generally lower in Experiment 4.

3.7.3 Discussion

Matute et al. (1996, 2002) proposed that the wording of the test question used to request participants' ratings is an important variable in modulating cue-interaction. Specifically, they argued that if participants were asked to judge the contiguity between a pair of cues and an outcome, or between a pair of effects and a cause, then the cues should not interact. On the other hand, if asked to judge the causality between a pair of causes and an effect, then the causes should interact. The results from Experiment 2 indicate that participants' sensitivity to causal asymmetry was not influenced by whether they were asked a causal or contiguous test question. Therefore, the second alleged causal model influence of the causal versus contiguous wording of the test question was ineffective.

3.8 Experiment 3: Integration

In discussing the results from their four experiments, Tangen and Allan (submitted) drew an analogy between their findings and the primacy-recency effect obtained by Collins and

Shanks (in press). In the primacy-recency effect, Collins and Shanks described the "momentary" strategy where judgments reflect the current associative strength of the cue, and the "integrative" strategy where participants do not constrain their judgments on the current perception of the relationship, but instead integrate information across a number of trials. Tangen and Allan proposed that similar strategies may be operating in assessing causally asymmetric events. Participants' sensitivity to causal structure can vary in degree according to what they are asked about the events and according to their experience with them. They argued that if participants were asked to use the causal model for some particular purpose, then they would likely be more sensitive to the causal structure. Similarly, if participants were asked an integrative rating question that reflected the general relationship of the events, then their assessments would better coincide with the predictions made by causal-model theory.

Experiment 3 was designed to investigate whether participants' ratings would better reflect the causal structure of the events if asked to provide an integrative rating, i.e., in contrast to the prediction responses required on every trial, and the four ratings requested every 16 trials. Participants were told that a pharmaceutical company was interested in manufacturing the two chemicals (2C-1E) or the bacterial strain (2E-1C) based on their observations. They were asked to assess the effectiveness of each chemical on the survival of the bacteria (2C-1E) or assess the effectiveness of the bacterial strain on the production of each chemical (2E-1C) based on the entire set of 64 trials. The wording of this "integration" question was intended to suggest that they should consider the causal description of the events rather than identifying them as generic cues and outcomes.

3.8.1 Method

The design of Experiment 3 was similar to the previous two experiments. Participants were asked to provide prediction responses on each trial and rate the relationship every 16 trials based on the contingencies presented in Table 3.1. After rating each chemical following Trial 64, participants were presented with the following dialogue:

2C-1E: You have just observed a series of 64 experimental trials. DrugCorp - a large pharmaceutical company - is interested in your assessment of each chemical on the survival of the bacteria and they intend to allocate funds to manufacture the chemicals on the basis of your observations. It is important that your estimate best reflects the effectiveness of each chemical across the entire set of 64 trials. Please use the scrollbar below each chemical to make a rating on a scale from -100 to +100.

2E-1C: You have just observed a series of 64 experimental trials. DrugCorp - a large pharmaceutical company - is interested in your assessment of the bacteria on the production of each chemical and they intend to allocate funds to manufacture the bacteria on the basis of your observations. It is important that your estimate best reflects the effectiveness of the bacteria across the entire set of 64 trials. Please use the scrollbar below each chemical to make a rating on a scale from -100 to +100.

After providing a rating for each chemical, another series of trials would begin based on one of the four contingency pairs listed in Table 3.1.

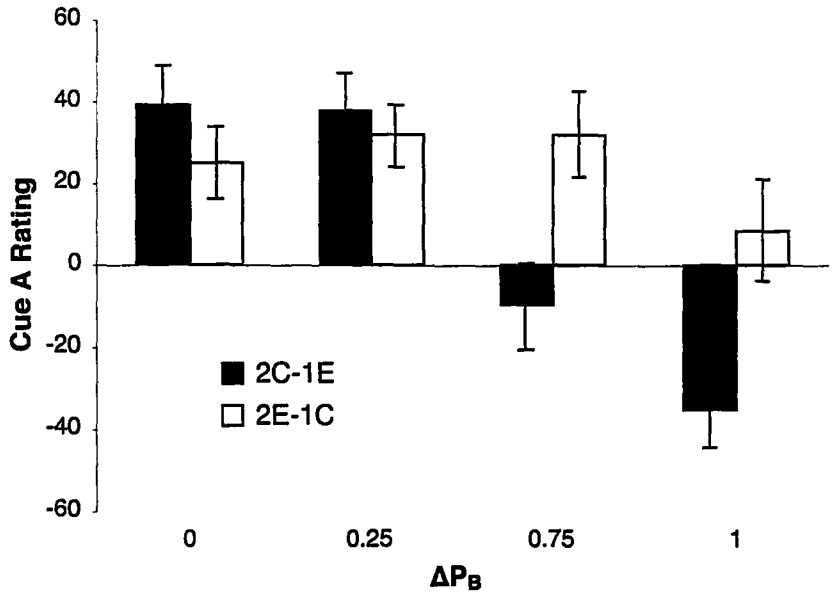

Figure 3.5: Mean ratings in Experiment 3 after 64 trials of Cue A (Figure 3.5). The ratings are shown as a function of ΔP_B (0, 0.25, 0.75, 1) separately for each of the two conditions. Error-bars represent standard errors of the means.

3.8.2 Results

In Experiment 3, there were three dependent measures, ratings, integrative ratings, and predictions.

Ratings

Figure 3.5 illustrates the mean ratings of Cue A after 64 trials and Table 3.5 provides the mean ratings of Cue A after 32, 48, and 64 trials. The data are presented for each of the four contingency pairs. The rating data from Experiment 3 seem to deviate slightly relative to the previous two experiments and the results obtained in Tangen and Allan (submitted, Experiment 4). Cue-interaction seems to be less evident in the overall ratings for the 2E-1C scenario than in the previous two experiments.

As in the previous experiments, the data from Tangen and Allan (submitted, Experiment 4) was used as a template, and a 2 (experiment: Exp 4, Exp 3) × 2 (scenario: 2C-1E, 2E-1C) × 4 (ΔP_B: 0, 0.25, 0.75, 1) × 3 (trial: 32, 48, 64) mixed ANOVA was conducted. Unlike the previous experiments, there was a significant effect of the causal model as indicated by the experiment × scenario × ΔP_B interaction, $F(3, 228) = 3.04$, $p < .05$. The only other effect (that interacted with experiment) to reach significance was a four-way interaction between experiment, scenario, ΔP_B, and trial, $F(6, 456) = 2.65$, $p < .05$ indicating that the

Table 3.5: Experiment 3 mean ratings and estimated ΔP values of Cue A conditional on the presence of Cue B ($est\Delta P_{A|B}$) and the absence of Cue B ($est\Delta P_{A|\sim B}$) after 32, 48, and 64 trials.

		2C-1E							
		0.5/0		0.5/0.25		0.5/0.75		0.5/1	
32	Rating	31.4	(8)	42.1	(9.8)	6.2	(10.7)	-13.3	(9.3)
	$est\Delta P_{A\|B}$	0.44	(0.07)	0.45	(0.08)	0.26	(0.06)	0.12	(0.06)
	$est\Delta P_{A\|\sim B}$	0.52	(0.07)	0.56	(0.06)	0.33	(0.07)	0.32	(0.08)
48	Rating	39.9	(9.1)	36.9	(8.7)	-4.6	(10.6)	-23.6	(8.9)
	$est\Delta P_{A\|B}$	0.52	(0.06)	0.44	(0.06)	0.17	(0.04)	0.05	(0.03)
	$est\Delta P_{A\|\sim B}$	0.58	(0.06)	0.57	(0.05)	0.27	(0.06)	0.21	(0.07)
64	Rating	39.1	(10)	37.8	(8.9)	-9.8	(10.5)	-35.3	(9.2)
	$est\Delta P_{A\|B}$	0.53	(0.07)	0.42	(0.06)	0.1	(0.04)	0.04	(0.03)
	$est\Delta P_{A\|\sim B}$	0.59	(0.06)	0.56	(0.05)	0.27	(0.06)	0.13	(0.04)
		2E-1C							
		0.5/0		0.5/0.25		0.5/0.75		0.5/1	
32	Rating	27.2	(10.2)	19.7	(11)	50.9	(9.3)	4	(13.2)
	$est\Delta P_{A\|B}$	0.29	(0.07)	0.36	(0.07)	0.16	(0.04)	0.18	(0.07)
	$est\Delta P_{A\|\sim B}$	0.51	(0.08)	0.48	(0.06)	0.44	(0.08)	0.21	(0.06)
48	Rating	29.7	(7.8)	24.3	(10.8)	34	(9.7)	-1.4	(11.6)
	$est\Delta P_{A\|B}$	0.36	(0.07)	0.38	(0.06)	0.17	(0.06)	0.14	(0.05)
	$est\Delta P_{A\|\sim B}$	0.55	(0.07)	0.47	(0.05)	0.47	(0.07)	0.2	(0.06)
64	Rating	24.9	(8.7)	31.6	(7.5)	31.9	(10.6)	8.5	(12.4)
	$est\Delta P_{A\|B}$	0.4	(0.07)	0.42	(0.06)	0.17	(0.06)	0.12	(0.05)
	$est\Delta P_{A\|\sim B}$	0.58	(0.06)	0.53	(0.06)	0.45	(0.06)	0.19	(0.05)

Note. Standard errors of the means are given in parentheses.

Table 3.6: Experiment 3 integrative ratings of Cue A.

	2C-1E			
	0.5/0	0.5/0.25	0.5/0.75	0.5/1
Rating	41.3	33.6	-4.4	-27.9
	(7.9)	(8.8)	(9.3)	(10)
	2E-1C			
	0.5/0	0.5/0.25	0.5/0.75	0.5/1
Rating	37.8	46.4	45.1	-3.9
	(7.1)	(8)	(8.9)	(11.6)

Note. Standard errors of the means are given in parentheses.

influence of the causal model did not dissipate over trials in the present experiment as it did in Experiment 4.

Integrative Ratings

Table 3.6 presents the mean integrative ratings of Cue A for each of the four contingency pairs after 64 trials. As in the ratings data, there seems to be an effect of the causal model. A 2 (scenario: 2C-1E, 2E-1C) × 4 (ΔP_B: 0, 0.25, 0.75, 1) mixed ANOVA conducted on the integrative judgments of Cue A confirms these observations. The ANOVA revealed significant main effects of scenario, $F(1, 38) = 6.66$, $p < .05$, and ΔP_B, $F(3, 114) = 20.79$, $p < .001$, as well as a significant interaction between them, $F(3, 114) = 3.73$, $p < .05$.

Predictions

Table 3.5 also depicts the estimated ΔP values conditional on the presence and absence of Cue B. The data are presented for each of the four contingency pairs after 32, 48, and 64 trials. The prediction response data do not seem to differ from those in the previous two experiments. Again, using the data from Tangen and Allan (submitted, Experiment 4) as a template, a 2 (experiment: Exp 4, Exp 3) × 2 (scenario: 2C-1E, 2E-1C) × 4 (ΔP_B: 0, 0.25, 0.75, 1) × 3 (trial: 32, 48, 64) × 2 (Cue B status: present, absent) mixed ANOVA was conducted on the estimated conditional ΔP values for Cue A. Unlike the rating data, there was no significant influence of the causal model as indicated by the experiment × scenario × ΔP_B interaction which was not significant, $F(3, 228) = 1.5$, $p > .05$. The only effect (that interacted with experiment) to reach significance in the prediction data was a three-way interaction between experiment, trial, and Cue B status, $F(2, 152) = 3.3$, $p < .05$.

3.8.3 Discussion

Overall, we see that participants are more sensitive to the structure of the causal relationship when asked to provide an integrative judgment than in the previous two experiments and in Tangen and Allan (submitted, Experiment 4). Therefore, the third alleged causal model influence of an integrative test question was effective. However, the dissociation between ratings and trial-by-trial predictions remains. Participants conditionalise in their prediction responses despite the circumstances. The effect of the integration question on ratings is not large by any means, but it nicely demonstrates how the relative weighting of causal and associative factors vary in a matter of degree depending on what is being asked about the events.

3.9 General Discussion

Tangen and Allan (submitted) have demonstrated that the contribution of causal and associative processes depends on what the participant is being asked about the events, and on their experience with those events. They presented two experiments providing evidence for high-level (causal reasoning) processes, and two experiments providing evidence for low-level (associative) processes. They argued that both factors influence causal assessment. The present series of experiments provides additional evidence for this dual-process argument.

Experiment 1 was designed to investigate whether emphasizing the direction of the causal relationship would influence assessments of that relationship as suggested by Waldmann and Holyoak (1997). Participants were presented with causal model "prompts" every 8 trials designed to remind them about the direction of the causal relationship. Even though they were presented with information signalling the direction of the relationship *immediately prior to rating that relationship*, it did not influence their sensitivity to causal directionality. Experiment 2 investigated the claim by Matute et al. (1996) that the wording of the test question used to request participants' ratings is an important variable in modulating participants' sensitivity to causal structure. Participants were presented either with a causal or contiguous test question prior to rating the contingency between the events. The data suggest that the wording of the test question results in cue-interaction for both scenarios. Following a prediction made by Tangen and Allan (submitted), Experiment 3 required participants to respond to an integrative test question designed so they would weight the causal description of the events more heavily than their associative nature by regarding the causal description of the events rather than identifying them as generic cues and outcomes. The manipulation was successful, in that the influence of the causal model was evident both in participants' ratings and integrative ratings.

According to causal-model theory, our knowledge of causal asymmetry provides us with the capacity to ignore the order in which events are presented thereby transforming them into causal-model representations that reflect their asymmetry (Waldmann, 2000). However, the extent and circumstances of our disregard for temporal order has not been specified. A conditional ΔP account of causal-model theory suggests that participants will conditionalise between causes, but not effects. Thus, causal assessments should coincide with conditional ΔP

when two causes produce a common effect (2C-1E), and unconditional ΔP when two effects result from a common cause (2E-1C). The pattern of results obtained in Experiments 1-3 are not consistent with causal-model theory in the strictest sense. Causal assessments of the moderately contingent Cue A ($\Delta P_A = 0.5$) should be identical in the 2E-1C scenario regardless of the contingency for Cue B. This is not the case. We see that cues interact in both scenarios. However, according to an associative account, the two cues should interact the same amount regardless of their causal description. This is not the case. We see a larger cue-interaction effect for the 2C-1E scenario than for 2E-1C scenario. We also see from Experiment 3 that this sensitivity to causal asymmetry can change depending on what participants are asked about the events.

The results from Experiments 1-3 indicate that causal assessments are influenced both by the causal description and the associative nature of the events. Throughout the three experiments, participants' A ratings seem to be higher in the 2E-1C scenario when it is paired with a stronger Cue B (indicating their sensitivity to the causal structure of the events). This difference is especially pronounced in both the rating and integration data from Experiment 3 when participants are provided with an integrative question. On the other hand, the analyses reveal significant cue-interaction effects for both scenarios in Experiments 1 and 2. As well, participants' prediction response data indicate no sensitivity whatsoever to causal asymmetry (indicating their sensitivity to the associative structure of the events). This mixture of high and low-level processing is consistent with the results from Tangen and Allan (submitted) and Price and Yates (1995) who argue for the joint contribution of associative and causal factors in judgments of contingency.

Furthermore, as predicted by Tangen and Allan (submitted), the balance of sensitivity to the associative and causal structure of the events can be shifted depending on what participants are asked about the events as indicated by the results from Experiment 3. Indeed, it seems as though assessing (in)sensitivity to causal asymmetry cannot be explained solely by an associative account or causal-model theory, but may be better explained in terms of their joint contribution.

Chapter 4

The Relative Effect of Cue-interaction

> Having given the number of instances respectively in which things are both thus and so, in which they are thus but not so, in which they are so but not thus, and in which they are neither thus nor so, it is required to eliminate the general quantitative relativity inhering in the mere thingness of the things, and to determine the special quantitative relativity subsisting between the thusness and the soness of the things.
>
> *Doolittle (1888)*

4.1 Preface

This chapter is reproduced from Tangen and Allan (2003) that was submitted to the *Quarterly Journal of Experimental Psychology. B, Comparative and Physiological Psychology* on June 4, 2002 and accepted for publication on September 3, 2002. The data presented in this paper represents a section of my research designed to examine the role of conditionalisation on judgments of contingency. In particular, we investigate how human reasoners evaluate the causal strength of a given cause while controlling for alternative causes. The experiments were designed to test several predictions derived from Spellman's conditional ΔP account. For example, the model predicts that under certain conditions, judgments of the effectiveness of a moderately positive predictor of an outcome will not be attenuated when training includes a more valid positive predictor. Participants, therefore, tend to judge the effectiveness of each cue relative to the other cue(s) that are presented simultaneously. The results are consistent with the conditional ΔP account. A competing account of cue-interaction is provided by the Rescorla-Wagner (RW) model. We present an algebraic derivation of the predictions of the RW model for the conditions specified by Spellman, and show that at asymptote the predictions of the RW model are identical to those of the conditional ΔP account.

4.2 Abstract

It is well established that two predictor cues (A and B) of a common outcome interact in that the judgment of the relationship between each cue and the outcome is influenced by the pairing history of the other cue with the common outcome. For example, when the contingency of A with the common outcome is weaker than the contingency of B with the common outcome, the rating of the predictiveness of A is reduced relative to a situation where only A is paired with the outcome. One explanation of such cue interaction effects is provided by the conditional ΔP account. Spellman (1996b) derived a counterintuitive prediction of the conditional ΔP account where cue interaction should not occur under certain conditions even though a relatively poor predictor of an outcome is paired with a relatively good predictor of that outcome. However, Spellman (1996b) did not provide data to evaluate this prediction. In the present paper, we report the relevant data and show that they are consistent with the conditional ΔP account. A competing account of cue interaction is provided by the Rescorla-Wagner (RW) model. We derive the predictions of the RW model for the conditions specified by Spellman (1996b), and show that at asymptote the predictions of the RW model are identical to those of the conditional ΔP account.

4.3 Introduction

While everyone agrees that organisms must be able to appreciate the relationships among events in their environment in order to survive, there is less agreement about the mechanism that detects such relationships (see Allan, 1993; Shanks, 1993; Shanks, Holyoak, & Medin, 1996). Figure 4.1 presents the standard 2×2 contingency matrix for the generic laboratory task used to study how human observers make judgements about the relationship between two binary variables (see Allan, 1980). In such tasks, a cue is either present (A) or absent (∼A) and the outcome is either present (O) or absent (∼O)[1]. The four letters in the cells (a, b, c, d) represent the joint frequency of occurrence of the four possible cue-outcome combinations. After a series of trials on which each of the four cue-outcome combinations are presented with a pre-defined frequency, the observer is asked about the relationship between the cue and the outcome. Many statistical rules have been proposed to describe the manner in which observers combine the information in Figure 4.1 to inform their judgement about the relationship between the cue and the outcome (see Allan, 1993). One such rule is the covariation or contingency between the cue and the outcome, ΔP, which is the difference between two independent conditional probabilities (Allan, 1980). Referring to Figure 4.1,

$$\Delta P = P(O|A) - P(O| \sim A) = \frac{a}{a+b} - \frac{c}{c+d} \tag{4.1}$$

Figure 4.1 conceptualizes the judgment task for one cue and one outcome. Recent research has been concerned with judgments in situations involving multiple cues and a common outcome. The data from such experiments indicate that multiple cues often interact in that the

[1]We use "cue" and "outcome" to refer to antecedent and subsequent event(s) respectively without any causal connotation.

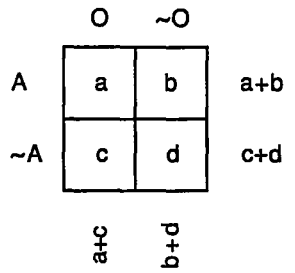

Figure 4.1: Summary 2×2 contingency matrix illustrating each of the possible cause-effect combinations for one cue and one outcome. Each cell represents the frequency of each event type.

judgment of the relationship between each cue and the common outcome is influenced by the pairing history of the other cues with the outcome. For example, consider two cues, A and B, where the contingency of A with the outcome is weaker than the contingency of B with the outcome. The usual finding is that the rating of the relationship of A with the outcome is reduced relative to a situation where only A is paired with the outcome. Depending on the task used and the theoretical framework of the experiment, this reduction in the rating of A has been referred to as discounting, blocking, or overshadowing. As many have reported (see Allan, 1993), ΔP as defined in Equation 4.1 does not account for such cue interaction data.

In fact, the calculation of ΔP for A in Equation 4.1 ignores the presence of alternative cues in that $P(O|A)$ is based on all A trials and $P(O| \sim A)$ is based on all ~A trials. A number of investigators (e.g., Cheng & Novick, 1990, 1992; Spellman, 1996a, 1996b) have argued that the ΔP rule should be applied across a "focal set" of trial types rather than across all trials:

$$\Delta P = P(O|AC) - P(O| \sim AC) \tag{4.2}$$

where C represents all cues that are constant on A and ~A trials. That is, focal-set ΔP is based on trials that are identical except for the presence or the absence of A. For two cues, A and B, the determination of ΔP for A should be conditional on the presence of B ($\Delta P_{A|B}$) and on the absence of B $\Delta P_{A|\sim B}$, where

$$\Delta P_{A|B} = P(O|AB) - P(O| \sim AB) \tag{4.3}$$

$$\Delta P_{A|\sim B} = P(O|A \sim B) - P(O| \sim A \sim B) \tag{4.4}$$

Spellman (1996b) derived a number of interesting, and often counterintuitive, predictions from the conditional ΔP account of cue-interaction. One purpose of the present chapter is to provide the data to evaluate one of these predictions.

Studies concerned with a conditional ΔP account of cue-interaction have often used the single-phase simultaneous blocking task introduced into the human contingency judgement literature by Baker et al. (1993, see also, Price and Yates, 1993, 1995). In this task, there are

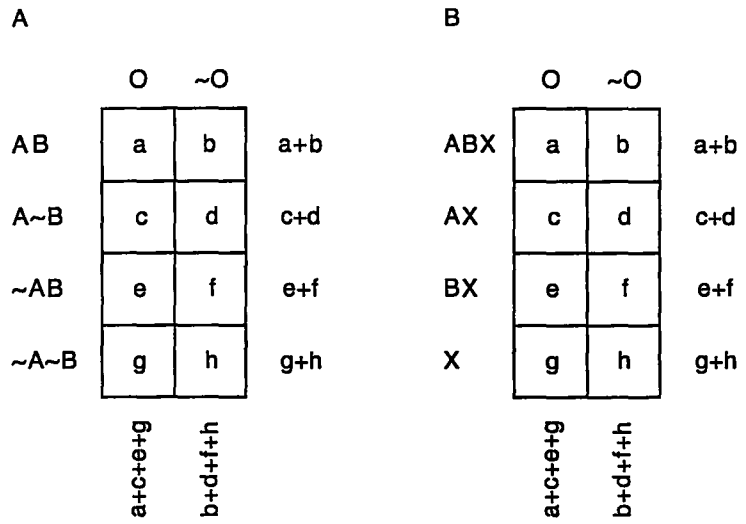

Figure 4.2: Summary 4×2 contingency matrix for the single-phase simultaneous blocking task. The notation in Figure 4.2a is for the conditional ΔP account, and the notation in Figure 4.2b is for the Rescorla-Wagner model.

two cues, A and B, followed by a common outcome (see Figure 4.2a). One of four possible cue combinations occurs on each trial: both cues present (AB), one cue present and the other absent (A~B or ~AB), and both cues absent (~A~B). For each cue combination, the outcome either occurs (O) or does not occur (~O). After a series of trials, where each of the eight cue-outcome combinations is presented with a pre-defined frequency, the observer is asked about the relationship between each cue and the outcome. The usual finding in such studies is that observers rate A as less predictive of the outcome when B is a strong positive predictor of the outcome than when B is a poor predictor (e.g., Baker et al., 1993).[2] That is, ratings of the predictive value of A depend on the predictive value of B.

Table 4.1 presents example trial frequencies that result in a moderately positive contingency for A ($\Delta P_A = .5$) in the presence of two different contingencies for B: a perfect positive contingency ($\Delta P_B = 1$) and a zero contingency ($\Delta P_B = 0$). We will use the notation introduced by Baker et al. (1993) to represent the unconditional contingencies of the two cues, $\Delta P_A/\Delta P_B$. Thus, the designation for the two examples in Table 4.1 is .5/1 and .5/0, where the first number represents ΔP_A and the second number represents ΔP_B. Table 4.1 also provides the two conditional ΔP values for A ($\Delta P_{A|B}$ and $\Delta P_{A|\sim B}$) and the two conditional ΔP

[2]While typically the rating of a moderately positive predictor decreases when it is paired with a stronger positive predictor, there is evidence that the rating of a moderately positive predictor increases when it is paired with a stronger negative predictor (e.g., Baker et al., 1993; Vallée-Tourangeau, Murphy, & Baker, 1998). In the present experiment, only positive unconditional ΔP values are used.

Table 4.1: Frequency matrices, and unconditional and conditional probabilities for the situation where a moderately predictive cue is paired with a perfect predictor (.5/1) and for the situation where a moderately predictive cue is paired with a non predictive cue (.5/0).

Trial Type	.5/0	.5/1
ABO	6	12
A∼BO	6	0
∼ABO	2	4
∼A∼BO	2	0
AB∼O	2	0
A∼B∼O	2	4
∼AB∼O	6	0
∼A∼B∼O	6	12
# of Trials	32	32
ΔP_A	.5	.5
$\Delta P_{A\|B}$	.5	0
$\Delta P_{A\|\sim B}$	.5	0
ΔP_B	0	1
$\Delta P_{B\|A}$	0	1
$\Delta P_{B\|\sim A}$	0	1

values for B ($\Delta P_{B|A}$ and $\Delta P_{B|\sim A}$). For the 4 x 2 matrix,

$$\Delta P_A = P(O|A) - P(O|\sim A) = \frac{a+c}{a+b+c+d} - \frac{e+g}{e+f+g+h} \tag{4.5}$$

$$\Delta P_B = P(O|B) - P(O|\sim B) = \frac{a+e}{a+b+e+f} - \frac{c+g}{c+d+g+h} \tag{4.6}$$

$$\Delta P_{A|B} = P(O|AB) - P(O|\sim AB) = \frac{a}{a+b} - \frac{e}{e+f} \tag{4.7}$$

$$\Delta P_{A|\sim B} = P(O|A\sim B) - P(O|\sim A\sim B) = \frac{c}{c+d} - \frac{g}{g+h} \tag{4.8}$$

$$\Delta P_{B|A} = P(O|BA) - P(O|\sim BA) = \frac{a}{a+b} - \frac{c}{c+d} \tag{4.9}$$

$$\Delta P_{B|\sim A} = P(O|B\sim A) - P(O|\sim B\sim A) = \frac{e}{e+f} - \frac{g}{g+h} \tag{4.10}$$

When $\Delta P_B = 0$ in Table 4.1, the unconditional and conditional ΔP values for A are the same; i.e., $\Delta P_A = \Delta P_{A|B} = \Delta P_{A|\sim B} = .5$. However, when $\Delta P_B = 1$, this is not the case; whereas $\Delta P_A = .5$, $\Delta P_{A|B} = \Delta P_{A|\sim B} = 0$. Thus, if observers are basing their judgements of A conditional on B, they should indeed judge A as less predictive of the outcome when $\Delta P_B = 1$ and $\Delta P_{A|B} = \Delta P_{A|\sim B} = 0$, then when $\Delta P_B = 0$ and $\Delta P_{A|B} = \Delta P_{A|\sim B} = .5$.

Spellman (1996b) showed that if the two conditional contingencies for A are equal, then the two conditional contingencies for B are also equal.

$$\text{if } \Delta P_{A|B} = \Delta P_{A|\sim B}, \text{ then } \Delta P_{B|A} = \Delta P_{B|\sim A} \qquad \text{(Property 1)}$$

She also showed that when Property 1 held and when the four row frequencies of the 4×2 matrix in Figure 4.2a are identical [i.e., $(a+b) = (c+d) = (e+f) = (g+h)$], then for each cue the unconditional and conditional probabilities are the same.

$$\Delta P_A = \Delta P_{A|B} = \Delta P_{A|\sim B} \text{ and } \Delta P_B = \Delta P_{B|A} = \Delta P_{B|\sim A} \qquad \text{(Property 4)}$$

Spellman (1996b, Figure 8) considered the three conditions illustrated in Table 4.2.[3] In all three conditions $\Delta P_A = .33$. In the two .33/.67 conditions, the row frequencies of the 4×2 matrix are the same in the Equal condition but are not the same in the Unequal condition. In the Equal condition, for each cue the conditional contingencies are the same as the unconditional contingency, whereas in the Unequal condition they are not. Thus, according to the conditional ΔP account, ratings in the two .33/.67 conditions should differ. Specifically,

[3]Spellman (1996b) labelled her three conditions baseline, no discounting and discounting. We prefer the non theoretical and empirical labels baseline, equal and unequal respectively. Also, it should be noted that the Spellman (1996b) matrices are arranged differently than ours are, and therefore the letters a, b, ..., h frequently represent different cells.

Table 4.2: 4×2 matricies showing the frequencies suggested by Spellman (1996b, Figure 8) for three pairings of two cues.

Trial Type	.33/0 Baseline	.33/.67 Equal	.33/.67 Unequal
ABO	3	9	12
A∼BO	3	3	0
∼ABO	0	6	3
∼A∼BO	0	0	3
AB∼O	6	0	3
A∼B∼O	6	6	3
∼AB∼O	9	3	0
∼A∼B∼O	9	9	12
# of Trials	36	36	36
ΔP_A	.33	.33	.33
$\Delta P_{A\|B}$	.33	.33	-.20
$\Delta P_{A\|\sim B}$	.33	.33	-.20
ΔP_B	0	.67	.67
$\Delta P_{B\|A}$	0	.67	.80
$\Delta P_{B\|\sim A}$	0	.67	.80

ratings of A should be the same in .33/.67 Equal as in .33/0 Baseline and should be higher than in .33/.67 Unequal. Also, ratings of B should be lower in .33/.67 Equal than in .33/.67 Unequal.[4] While Spellman (1996b) discussed these interesting predictions of the conditional ΔP account, she did not provide the data to evaluate the predictions. One purpose of the present chapter is to provide these data.

The ΔP rules (unconditional and conditional) stipulate that the frequency of the co-occurrence of cues and outcomes are mathematically transformed into probabilities and that judgements are based on an arithmetic comparison of these probabilities. An alternative view is provided by associative models originally developed to account for Pavlovian conditioning in animal learning experiments. These models postulate that judgements are determined by associative links formed between contiguously presented cues and outcomes. One such associative model, proposed by Rescorla and Wagner (1972), has frequently been applied to human contingency judgements (see Allan, 1993). According to the Rescorla-Wagner model, a cue gains associative (predictive) strength only to the extent that it provides information about the occurrence of the outcome that is not available from another source. The change in the predictive strength of the outcome by the cue is proportional to the degree to which the outcome is unexpected or surprising given all the alternative cues present on that trial. More precisely, the predictive strength of A (V_A) will change on each trial that it is presented according to the standard linear operator equation

$$\Delta V = \alpha\beta\left(\lambda - \sum V\right), \tag{4.11}$$

where ΔV_A is the change in predictive strength of A, α and β are learning rate parameters that depend on the salience of the cue and the effectiveness of the outcome respectively, λ is the maximum amount of predictive strength supported by the outcome, and $\sum V$ is the sum of the predictive strengths of all cues present on that trial.

The essence of the Rescorla-Wagner model is cue competition: there is a limit (λ) to the amount of predictive strength that an outcome can support. This limited amount of predictive strength is allocated among all cues present on the trial. If one cue acquires predictive strength, then all other cues that are present at the same time must get less. In the Rescorla-Wagner model, a cue never occurs in isolation but is always compounded with contextual (background) cues. The 4×2 matrix for the single-phase blocking design, as conceptualized by the Rescorla-Wagner model, is shown in Figure 4.2b. In the Rescorla-Wagner model, the outcome is associated with the explicit cues (A and B) and also with non-explicit or contextual cues (X). These non-explicit cues compete with A and B for predictive strength.

Chapman and Robbins (1990) derived the relationship between ΔP for the 2×2 matrix (Figure 4.1) and V_A. They showed that at asymptote (i.e., $\Delta V_A = 0$), $V_A = \Delta P$ if β_O (outcome present) = $\beta_{\sim O}$ (outcome absent)[5]. That is, the predictive strength of A as described

[4]The predictions from the conditional ΔP account are ordinal in that only the relative rankings of conditions are specified.

[5]The assumption that β_O (outcome present) = $\beta_{\sim O}$ (outcome absent) is often made in the human judgement task. While there are situations where the βs need to be unequal in order for the Rescorla-Wagner model to fit the data (see Lober & Shanks, 2000), there are other situations where equal βs do the job (see Allan, 1993).

by the Rescorla-Wagner model is identical to the contingency between A and the outcome as described by ΔP. Thus, the Rescorla-Wagner model explains how a sensitivity to ΔP emerges from a process that does not explicitly calculate ΔP. Chapman and Robbins (1990) derived the relationship between asymptotic V and ΔP for one cue. Another purpose of this chapter is to derive the relationship for two cues of a common outcome under the conditions specified by Properties 1 and 4 of Spellman (1996b).

4.4 Method

4.4.1 Observers

The observers were 20 undergraduate students enrolled in Psychology courses at McMaster University who participated for course credit. They had not participated in other experiments concerned with contingency judgements.

4.4.2 Apparatus

Up to four observers at a time performed the experiment on Power Macintosh computers located in separate rooms. The experiment was programmed in MetaCard 2.3.1.

The cues were chemicals and the outcome was a bacterial strain. The presence of a cue was indicated by a coloured three-dimensional animation of a chemical spinning on its axis and the absence of a cue was indicated by a faded unmoving grayscale picture of the chemical. Similarly, the presence of the outcome was indicated by a coloured animation of moving bacteria and the absence of the outcome was indicated by a faded grayscale picture of stationary bacteria. There were eight chemicals and four bacterial strains. The names of the chemicals and bacteria were fictitious and were adapted from Mehta (2000), as were the instructions presented in Appendix B. The coloured moving images (indicating the presence of cues and outcomes) were accompanied by a name (e.g., Chorbine Present). In contrast, the faded nonmoving images (indicating the absence of cues and outcomes) were not accompanied by text.

4.4.3 Procedure

The instructions for the experiment were presented to the observer on the computer monitor (see Appendix B). In brief, the observer was told that scientists have recently discovered three strains of bacteria that exist in the mammalian digestive system. For each strain, the scientists were testing whether certain pairs of chemicals aid in, interfere with, or have no effect on a strain's survival. To do this, a strain of bacteria was first placed in culture (petri dishes). After that, one chemical, the other chemical, both chemicals, or neither chemical was added to the bacterial culture. The scientists then verified whether or not the bacterial sample survived. The observer was also shown the rating scale that they would use to rate the effectiveness of each chemical on the survival of the bacteria. After reading the instructions, the observer was shown a summary screen of the eight cue-outcome combinations. Eight

practice trials were then presented where each of the eight cue-outcome combinations was presented in random order.

There were three conditions: .33/0 Baseline, .33/.67 Equal, and .33/.67 Unequal. The frequencies of the cue-outcome combinations (see Table 4.2) were those suggested by Spellman (1996b) except we doubled the total number of trials to 72. The three conditions were presented in a random order to each observer. The observer initiated a condition by clicking on the "Begin" button on the computer screen. The cue combination was presented, with A on the left and B on the right. The observer predicted whether or not the bacteria survived by clicking the appropriate button on the screen, and was then shown whether or not the bacteria survived and was provided feedback (Correct, Incorrect) on their choice. The next trial was initiated by a mouse click on the "Next Trial" button. Observers rated how strongly each chemical (Chemical A followed by Chemical B) affected the survival of the bacterial strain after trial 18, 36, 54, and 72. The ratings were made on a horizontal scrollbar that ranged from -100 at the left to 100 on the right, and was anchored at 0 in the middle. Observers made their ratings by moving a scrollbar left and right with the mouse.

Each of the three conditions were clearly labeled as separate scientific experiments with different chemical and bacteria images and names. For each observer, two chemicals and one bacterial strain was randomly assigned to each of the three conditions. The remaining two chemicals and the remaining bacterial strain was used in the eight practice trials.

4.5 Results

The order of presentation of the eight cue-outcome combinations was randomized over the entire block of 72 trials. To determine whether consistent values of ΔP were generated at each of the four rating points, conditional ΔP values for A ($\Delta P_{A|B}$ and $\Delta P_{A|\sim B}$) and for B ($\Delta P_{B|A}$ and $\Delta P_{B|\sim A}$) were calculated for each observer at trials 18, 36, 54, and 72. The conditional ΔP values were cumulative in that the 18 trial value was based on the first 18 trials, the 36 trial value was based on the first 36 trials, the 54 trial value was based on the first 54 trials and the 72 trial value was based on all the trials. The determination of ΔP was cumulative because with only 18 trials, row frequencies of zero occasionally occurred and ΔP values could not be calculated. With 36 (or more) trials, all row frequencies for all observers were greater than zero. Also, the instructions were designed to encourage observers to base their ratings on all of the previous trials, not just the previous block of trials. To ensure that the conditional ΔP values did not change as additional trials were added, a 3 (condition: .33/0 Baseline, .33/.67 Equal, and .33/.67 Unequal) $\times$ 3 (trial: 36, 54, 72) within-subject ANOVAs was conducted on each of the four conditional ΔP values. For each of the four ANOVAs the trial factor was not significant ($ps > .05$).

4.5.1 Rating Data

Mean ratings for each condition are plotted as a function of trial (36, 54, and 72) in Figure 4.3. Figure 4.3a presents the ratings for Cue A and Figure 4.3b presents the ratings for Cue B. According to the conditional ΔP account, ratings in the two .33/.67 conditions (Equal and

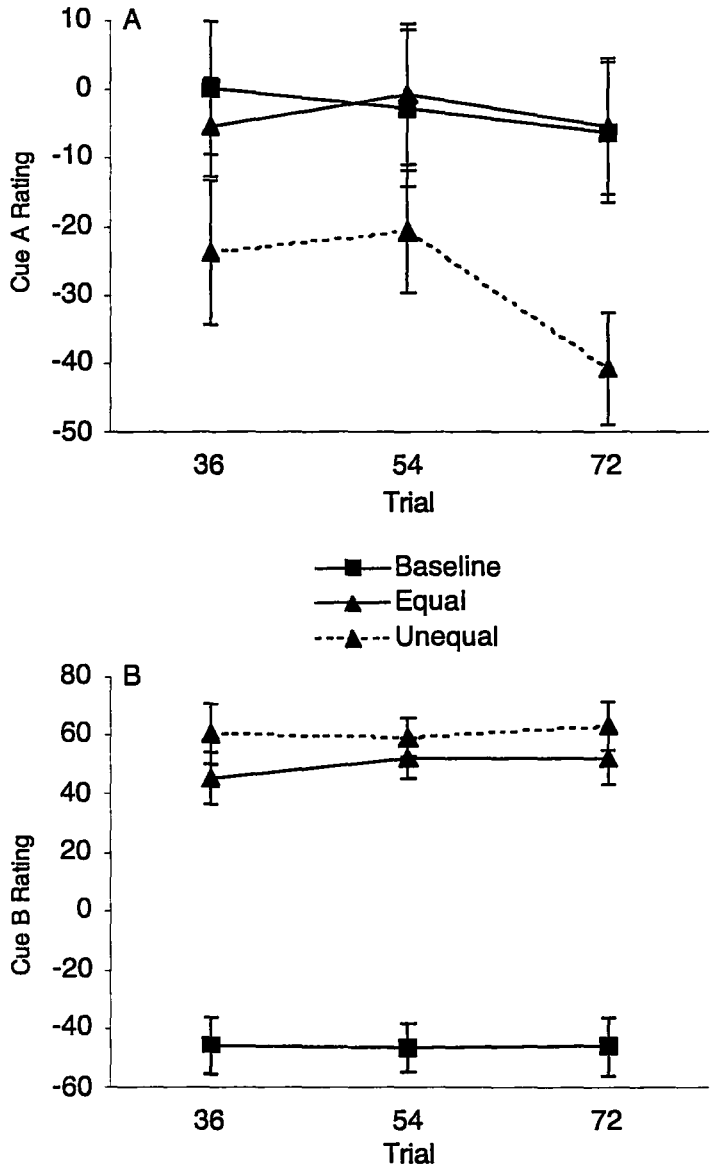

Figure 4.3: Mean ratings for each condition (Baseline, Equal, and Unequal) are plotted as a function of trial (36, 54, and 72). Figure 4.3a presents the data for Cue A and Figure 4.3b presents the data for Cue B.

Unequal) should differ. Ratings of A should be the same in .33/.67 Equal as in .33/0 Baseline and should be higher than in .33/.67 Unequal. Ratings of B should be lowest in Baseline, and should be lower in Equal than in Unequal. The pattern of results illustrated in Figure 1 for both cues are consistent with the conditional ΔP account. Ratings for A are similar in the Baseline and Equal conditions and are noticeably lower in the Unequal condition. Ratings for B are lowest in Baseline and highest in .33/.67 Unequal.

A 3 (condition: .33/0 Baseline, .33/.67 Equal, and .33/.67 Unequal) × 3 (trial: 36, 54, 72) within-subject ANOVA conducted on the ratings of A revealed a significant main effect of condition, $F(2, 38) = 4.77$, $p < .02$. Consistent with the ΔP account, the Tukey test showed that the ratings in the Unequal condition were significantly lower than in the Baseline, $p < .03$, and Equal, $p < .04$, conditions, and that the ratings in the Baseline and Equal conditions did not differ, $p > .05$. The trial main effect was not significant, $F(2, 38) = 2.10$, $p > .05$, nor did it interact with condition, $F(4, 76) = .54$, $p > .05$, indicating that the A ratings did not change after 36 trials. A similar ANOVA, conducted on the ratings for B, also revealed a significant main effect of condition, $F(2, 38) = 59.13$, $p < .001$. Consistent with the ΔP account, the Tukey test showed that the ratings in the Baseline condition were significantly lower than in the Equal, $p < .001$, and Unequal, $p < .001$, conditions. Although the difference between the Equal and Unequal conditions was in the direction predicted by the ΔP account, the difference was not significant, $p > .05$. Again, the trial main effect was not significant, $F(2, 38) = .24$, $p > .05$, nor did it interact with condition, $F(4, 76) = .24$, $p > .05$, indicating that the B ratings did not change after 36 trials.

Overall, the ratings of the two cues were in accord with a conditional ΔP account. What is especially noteworthy is that the ΔP account predicts that there should be no cue-interaction when Properties 1 and 4 hold (condition .33/.67 Equal) and the data were in accord with this prediction. Observers' ratings of A were the same in the .33/.67 Equal condition as they were in the .33/0 Baseline condition.

4.5.2 Prediction Data

Ratings are not the only dependent measure in this experiment. On each trial, after seeing one of the four cue combinations, the observer predicted whether the bacterial strain would survive or not. These prediction responses can be used to provide estimates of conditional ΔP values (López et al., 1998). A 4×2 matrix, like the one shown in Figure 4.2a, can be constructed where the columns are the two prediction responses (Yes, No) rather than the actual outcomes. Figure 4.4 plots the estimated conditional ΔP values for A; $est\Delta P_{A|B}$ in Figure 4.4a and $est\Delta P_{A|\sim B}$ in Figure 4.4b. Whereas the programmed values of the two conditional ΔP values ($\Delta P_{A|B}$ and $\Delta P_{A|\sim B}$) were the same, Figure 4.4 suggests that estimated ΔP conditional on a present cue (Figure 4.4a) was lower than estimated ΔP conditional on an absent cue (Figure 4.4b); that is, $est\Delta P_{A|B} < est\Delta P_{A|\sim B}$.

A 2 (B status: present, absent) × 3 (condition: Baseline, Equal, Unequal) × 3 (trial: 36, 54, 72) within-subject ANOVA was conducted on the estimated conditional ΔP values for A. All three main effects were significant, and none of the interactions were significant. The significant main effect of B status, $F(1, 19) = 15.27$, $p < .001$, indicates that estimated

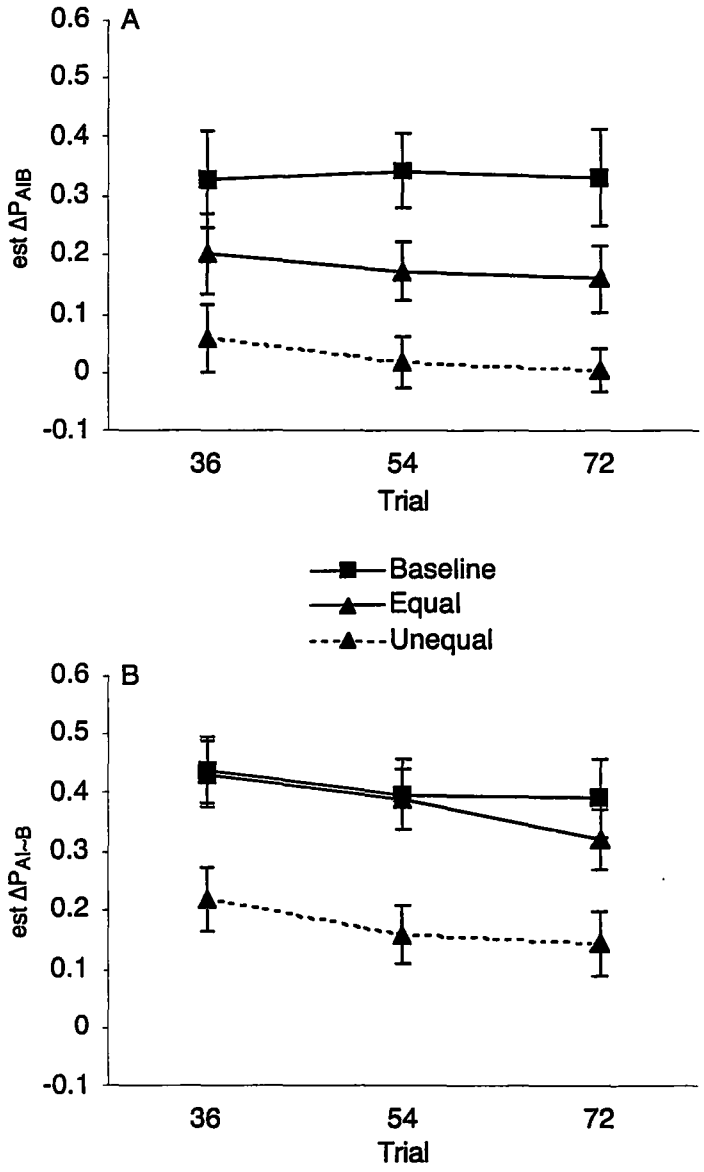

Figure 4.4: Mean estimated conditional ΔP values for Cue A are plotted as a function of trial (36, 54, and 72) for each condition (Baseline, Equal, and Unequal). $est\Delta P_{A|B}$ is shown in Figure 4.4a and $est\Delta P_{A|\sim B}$ is shown in Figure 4.4b.

conditional ΔP for A was lower when B was present (.18) than when it was absent (.32). The significant main effect of condition, $F(2, 38) = 11.18$, $p < .001$, was further analyzed using the Tukey test which showed that $est\Delta P$ in the Unequal condition (.10) was significantly lower than in the Baseline (.37), $p < .001$, and Equal (.28), $p < .01$, conditions, and that the Baseline and Equal conditions did not differ, $p > .05$. The Tukey test following the significant main effect of trial, $F(2, 38) = 6.22$, $p < .005$, indicated that estimated ΔP based on the first 36 trials (.28) was significantly higher than the value based all 72 trials (.23), $p < .004$.

Figure 4.5 plots the estimated conditional ΔP values for B; $est\Delta P_{B|A}$ in Figure 4.5a and $est\Delta P_{B|\sim A}$ in Figure 4.5b. As with Cue A, estimated ΔP conditional on a present cue (Figure 4.5a) was lower than estimated ΔP conditional on an absent cue (Figure 4.5b); that is, $est\Delta P_{B|A} < est\Delta P_{B|\sim A}$. A 2 (A status: present, absent) × 3 (condition: Baseline, Equal, Unequal) × 3 (trial: 36, 54, 72) within-subject ANOVA was conducted on the estimated conditional ΔP values for B. All three main effects were significant, as was the interaction between condition and trial. The significant main effect of A status, $F(1, 19) = 15.27$, $p < .001$, indicates that the estimated conditional ΔP value for B was lower when A was present (.33) than when it was absent (.47). The significant main effect of condition, $F(2, 38) = 53.46$, $p < .001$, was followed by the Tukey test which showed that $est\Delta P$ in the Baseline condition (-.01) was significantly lower than in the other two conditions, $ps < .001$, and that the Equal (.60) and Unequal (.61) conditions did not differ, $p > .05$. The Tukey test following the significant main effect of trial, $F(2, 38) = 21.08$, $p < .001$, indicated that estimated ΔP based on the first 36 trials (.35) was significantly lower than the value based on the first 54 trials (.42), $p < .001$ and the value based on all 72 trials (.43), $p < .001$. The Tukey test also revealed that the significant interaction between condition and trial, $F(4, 76) = 2.50$, $p < .05$, reflected that the trial effect was weakest in the Baseline condition and strongest in the Unequal condition.

In summary, the estimated conditional ΔP values provided a similar picture as the ratings for the three conditions. For Cue A, both estimated conditional ΔP and ratings were similar in the Baseline and Equal conditions and were lower in the Unequal condition. For Cue B, both estimated conditional ΔP and ratings were lowest in Baseline and highest in .33/.67 Unequal. The estimated conditional ΔP values differed from the ratings, however, in that they revealed trial effects for both cues. The greater variability in the ratings, in comparison to the estimated conditional ΔP values, might have masked the trial effects in the ratings.

Whereas the programmed values of the two conditional ΔP values for each cue did not differ, estimated ΔP conditional on a present cue was significantly lower than estimated ΔP conditional on an absent cue; $est\Delta P_{A|B} < est\Delta P_{A|\sim B}$, and $est\Delta P_{B|A} < est\Delta P_{B|\sim A}$. While many experiments reported in the literature have required prediction responses on each trial, usually these responses are not reported. For example, Spellman (1996a), Spellman, Price, and Logan (2001) asked observers to predict the outcome on each trial, but the prediction responses were not included as part of the data analysis. Our analysis of the prediction responses to provide estimates of conditional ΔP values indicates that these responses are informative. They provided confirming evidence for the prediction made by Spellman (1996b). They were less variable than the ratings, and showed trial effects that might have been masked in the ratings. Although the conditional ΔP account has not explicitly addressed the relationship of

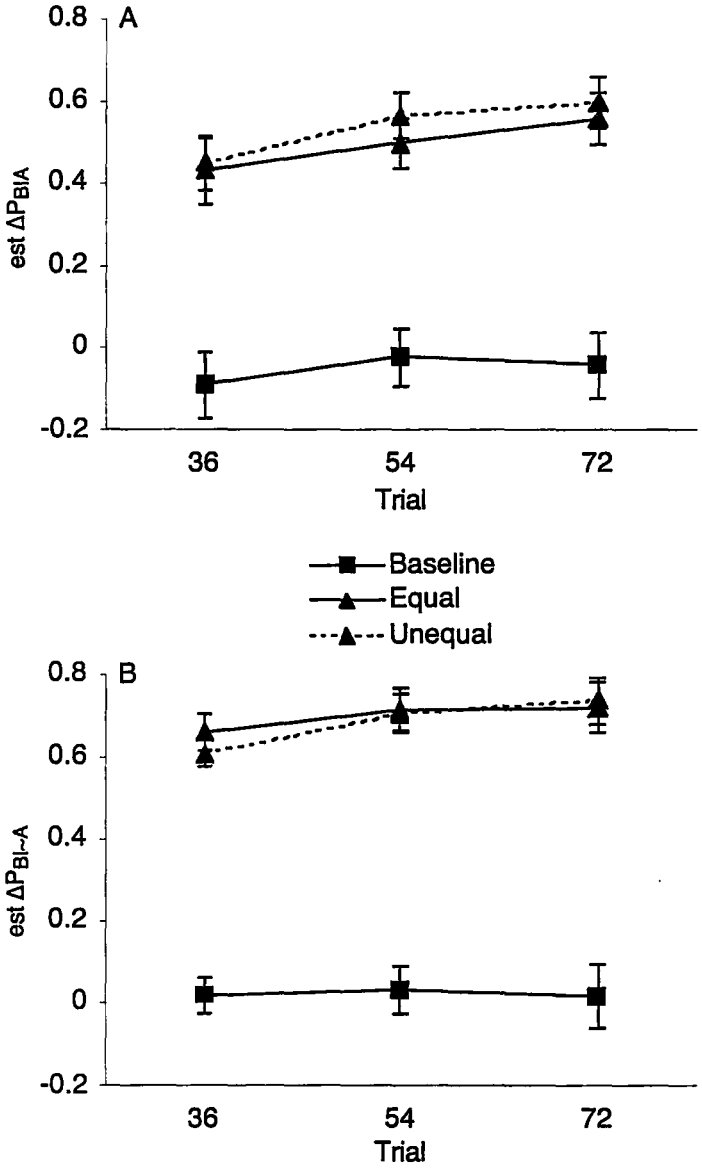

Figure 4.5: Mean estimated conditional ΔP values for Cue B are plotted as a function of trial (36, 54, and 72) for each condition (Baseline, Equal, and Unequal). $est\Delta P_{B|A}$ is shown in Figure 4.5a and $est\Delta P_{B|\sim A}$ is shown in Figure 4.5b.

estimated conditional ΔP to actual ΔP, one would expect them to be congruent. Our finding that the estimated values are not congruent with the actual values might be problematic for the conditional ΔP account.

4.5.3 The Rescorla-Wagner Model

We noted earlier that Chapman and Robbins (1990) derived the relationship between V_A and ΔP for the 2×2 matrix (Figure 4.1). They showed that at asymptote $\Delta V_A = 0$), $V_A = \Delta P$ if $\beta_O = \beta_{\sim O}$. In Appendix C, we show that at asymptote the Rescorla-Wagner model makes the same predictions as the conditional ΔP account for the 4×2 matrix when Properties 1 and 4 hold. Under these conditions, the Rescorla-Wagner model, like the ΔP account, predicts no cue-interaction. Consider again the 4×2 matrix in Figure 4.2. Spellman (1996b) showed that when Properties 1 and 4 hold,

$$\Delta P_{A|B} = \Delta P_{A|\sim B} = \Delta P_A = \frac{a-d}{a+b} \tag{4.12}$$

and

$$\Delta P_{B|A} = \Delta P_{B|\sim A} = \Delta P_B = \frac{a-c}{a+b} \tag{4.13}$$

In Appendix C, we show that when $\beta_O = \beta_{\sim O}$ and Properties 1 and 4 hold, at asymptote,

$$V_A = \frac{a-d}{a+b} = \Delta P_A \tag{4.14}$$

and

$$V_B = \frac{a-c}{a+b} = \Delta P_B \tag{4.15}$$

To our knowledge this prediction of the Rescorla-Wagner model has not been previously identified in the literature.[6] Rather it is has been assumed (albeit implicitly) that at asymptote there will be always be cue-interaction when a poor positive predictor (i.e., A) is paired with a better positive predictor (i.e., B). In the language of the Rescorla-Wagner model, B would block A. Appendix C identifies a set of conditions where this does not occur.

[6]Spellman (1996b, Footnote 8) noted that her simulations of the Rescorla-Wagner model resulted in asymptotic V values that were congruent with conditional ΔP values. These simulations, however, were not conducted for the frequency matrices used in the present experiment. Moreover, Spellman (1996b) did not provide an algebraic proof of the relationship between ΔP and V. Danks (in press) provides a detailed derivation of the behavior of the Rescorla-Wagner model, but he did not specifically address the predictions of the model for the situations examined in the present chapter. That is, he did not identify that under some situations the Rescorla-Wagner model predicts no cue-interaction.

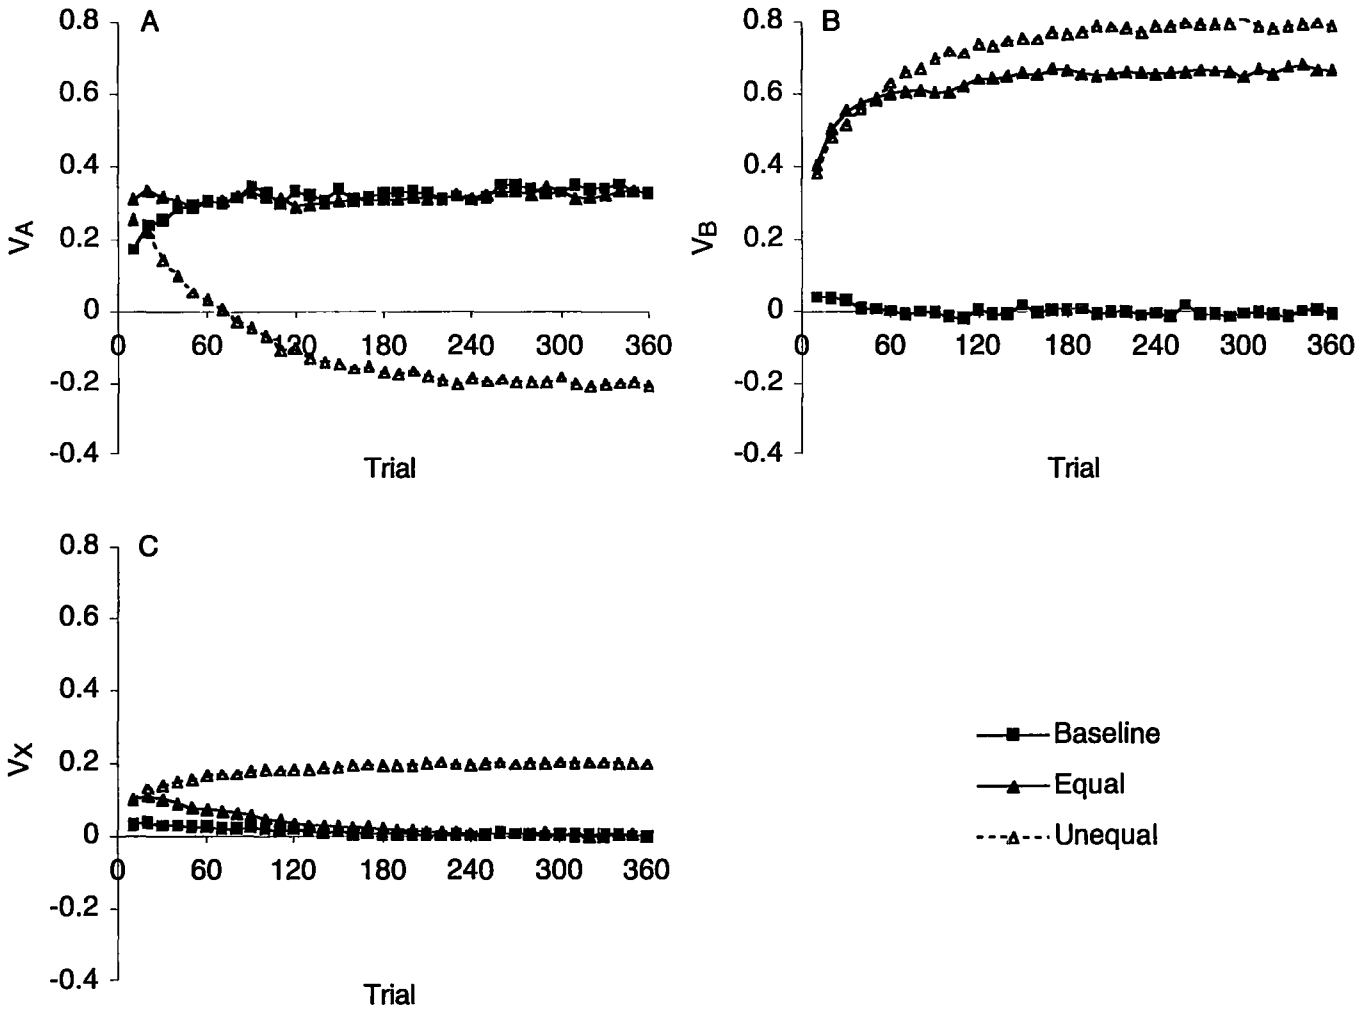

Figure 4.6: Rescorla-Wagner simulations for the three conditions (Baseline, Equal, and Unequal) are plotted in blocks of 10 trials. The predictive strength of Cue A (V_A) is shown in Figure 4.6a, of Cue B (V_B) in Figure 4.6b, and of the context (V_X) in Figure 4.6c.

Figure 4.6 presents Rescorla-Wagner simulations over 360 trials, in blocks of 10 trials, using an arbitrary set of parameter values for the three conditions used in the present experiment. The predictive strength of A (V_A) is shown in Figure 4.6a, of B (V_B) in Figure 4.6b, and of the context (V_X) in Figure 4.6c. The parameter values were those used by Allan (1993): $\alpha_A = \alpha_B = .9$, $\alpha_X = .2$, $\beta_O = \beta_{\sim O} = .2$, $\lambda_O = 1$, $\lambda_{\sim O} = 0$. Each simulation was based on 100 iterations. The simulations show that for the .33/.67 Equal condition the Rescorla-Wagner model predicts no cue-interaction at asymptote for A even though A is paired with a more predictive B.

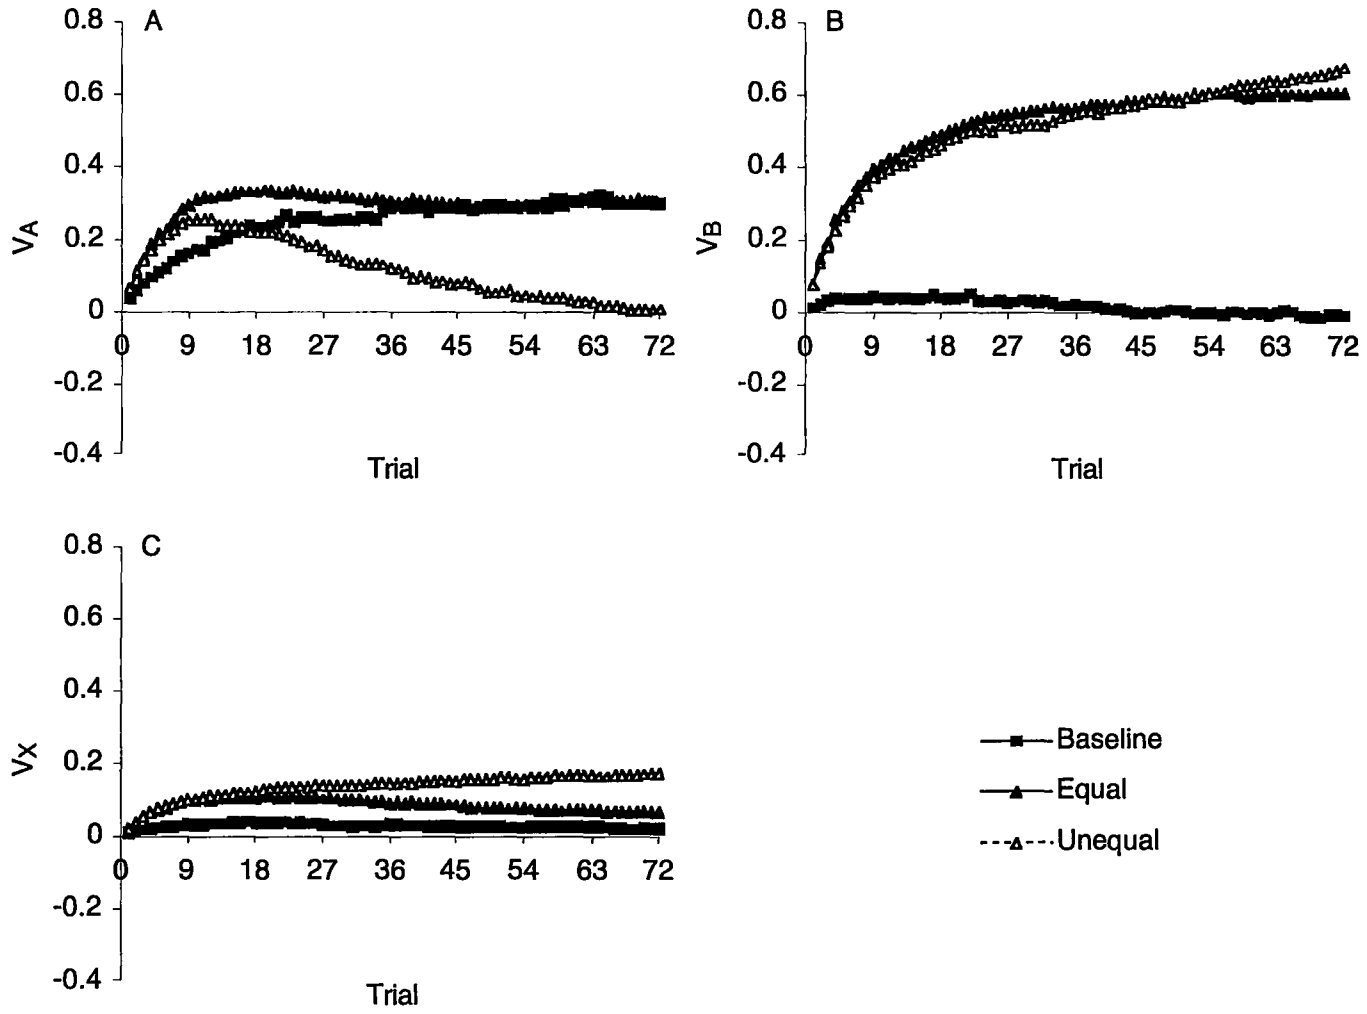

Figure 4.7: Rescorla-Wagner simulations for the three conditions (Baseline, Equal, and Unequal) are plotted for the first 72 trials. The predictive strength of Cue A (V_A) is shown in Figure 4.7a, of Cue B (V_B) in Figure 4.7b, and of the context (V_X) in Figure 4.7c.

While at asymptote the predictions of the Rescorla-Wagner model are identical to the predictions of the conditional ΔP account, Figure 4.6 illustrates that pre-asymptotically there are differences. Only the first 72 trials of the Rescorla-Wagner simulations are presented in Figure 4.7. On early trials, V_A is in fact the same for the two .33/.67 conditions, and it is higher than for the baseline condition. As the trials continue, V_A for the baseline condition increases to that of the Equal condition, and V_A for the Unequal condition gradually decreases. The course of predictive strength for A over trials is the result of the interaction of A not only with B but also with the context Cue X. In the Rescorla-Wagner model, the predictive strength of a cue only changes on trials that it is present. The context X has a special role, therefore, since it is present on all trials. Consider the two .33/.67 conditions. Early in training, V_X is similar for the two conditions. With extended training, V_X for the Unequal condition exceeds that for the Equal condition. Because A competes with X, as X gains predictive strength, A loses predictive strength.

Our earlier analysis of the rating data was conducted within the framework of the conditional ΔP account. Since it was not always possible to calculate a conditional ΔP value at 18 trials, we did not include the 18 trial rating in our statistical analyses. However, zero row frequencies are not problematic for the RW model since it is not based on conditional ΔPs. We reanalyzed the rating data for A and B including the 18 trial ratings. A 3 (condition: .33/0 Baseline, .33/.67 Equal, and .33/.67 Unequal) × 4 (trial: 18, 36, 54, 72) within-subject ANOVA conducted on the ratings of A revealed a significant main effect of trials, $F(3, 57) = 3.54$, $p < .02$. A similar ANOVA on the B ratings revealed a significant interaction of trials with conditions, $F(6, 114) = 2.86$, $p < .02$. Thus, while trial was not a significant factor when the 18 trial ratings were omitted, it was when they were included.

4.6 Discussion

Our data support the prediction made by Spellman (1996b) that under certain circumstances cue-interaction should not occur even though a relatively poor predictor is paired with a relatively good predictor. Spellman (1996b) showed that the conditional ΔP account predicts this outcome. In the present chapter we show that the Rescorla-Wagner model also predicts this outcome at asymptote.

While the conditional ΔP account predicts the main finding of the present experiment (the absence of cue-interaction), other aspects of the data are inconsistent with this account. Whereas the programmed values of the two conditional ΔP values for each cue did not differ, estimated ΔP conditional on a present cue was significantly lower than estimated ΔP conditional on an absent cue; $est\Delta P_{A|B} < est\Delta P_{A|\sim B}$, and $est\Delta P_{B|A} < est\Delta P_{B|\sim A}$. The conditional ΔP account, at least in its present form, does not predict these differences.

Trial effects were present in our data. Accounts based on ΔP, unconditional or conditional, are unable to incorporate trial effects. While estimates of conditional probabilities, and therefore ΔP, become more accurate with increasing sample size (trials), the mean estimate is independent of sample size.

The debate about the relative merits of a rule based model, such as conditional ΔP, or

an associative model, such as the Rescorla-Wagner model, has a long history (Allan, 1993; Shanks et al., 1996). Three major challenges for the unconditional ΔP account of contingency judgements (Equation 4.1) were cue-interaction effects, trial effects and density effects. The formulation of the conditional ΔP account was mainly in response to the cue-interaction challenge. While the unconditional ΔP account is unable to encompass cue-interaction effects, the conditional ΔP account has been more successful. The present data clearly show that at asymptote judgements are consistent with the predictions of the conditional ΔP account. Associative models were developed specifically to account for cue-interaction data generated in animal learning tasks. Overall, associative models have also been successful in accounting for the human cue-interaction data. There have been some interesting deviations of the human data from the predictions of associative models such as retrospective revaluation. These deviations have led to modifications of the associative models (e.g., Van Hamme & Wasserman, 1994; Dickinson & Burke, 1996).

Trial effects in contingency judgement tasks are ubiquitous. Many studies have shown that judgements change systematically over trials (see Allan, 1993). While trial effects are problematic for ΔP accounts, they are easily encompassed by associative models.

Many studies have shown that judgements are influenced not only by ΔP, but also by the frequency of cue present relative to the frequency of cue absent (a cue density effect) and by the frequency of outcome present relative to the frequency of outcome absent (an outcome density effect) (e.g., Allan & Jenkins, 1983). The Power PC model proposed by Cheng (1997) provides a conditional ΔP account which addresses density effects. However, a number of investigators have demonstrated that the Power PC model has a variety of difficulties and often does not provide an adequate account of the data (e.g., Allan, 2003; Lober & Shanks, 2000; Vallée-Tourangeau, Murphy, & Drew, 1997; Vallée-Tourangeau, Murphy, Drew, & Baker, 1998). Cue and outcome density effects are not problematic for associative models (see Allan, 1993).

In summary, the present chapter had two major purposes. One purpose was to provide data relevant to evaluating a prediction made by Spellman (1996b). Our data are consistent with the prediction made by Spellman (1996b). The second purpose was to show mathematically that under the conditions specified by Properties 1 and 4 made by Spellman (1996b), the Rescorla-Wagner model makes identical predictions at asymptote as the conditional ΔP account. In addition, our analysis of the prediction responses to provide estimates of conditional ΔP indicates that these responses, which are usually ignored, are informative. While our research was not designed to provide a strong test to differentiate a conditional ΔP account from an associative account, the dependence of the estimated values on the presence and absence of the other cue and the presence of trial effects are problematic for the conditional ΔP account.

4.7 Postscript

For each of the experiments in Chapters 2-4, the event frequencies were selected so the conditional ΔP values would be identical given the presence and absence of the alternate cue, where $\Delta P_{A|B} = \Delta P_{A|\sim B}$ and $\Delta P_{B|A} = \Delta P_{B|\sim A}$. However, as indicated by the trial-by-trial

prediction response data, the estimated ΔP values were *not* identical given the presence and absence of the alternate cue. That is, while actual $\Delta P_{A|B} = \Delta P_{A|\sim B}$ and $\Delta P_{B|A} = \Delta P_{B|\sim A}$, participants estimated $\Delta P_{A|B} \neq \Delta P_{A|\sim B}$ and $\Delta P_{B|A} \neq \Delta P_{B|\sim A}$. A conditional ΔP account does not predict these differences. How would participants rate a causal relationship in which the contingency conditional on the presence of an alternative cause is different than the contingency conditional on the absence of an alternative cause? And, what does the Rescorla-Wagner model predict under these circumstances?

4.8 Method

The participants were 56 undergraduate students enrolled in Psychology courses at McMaster University who participated for course credit. 22 participants participated in Experiment 1, 17 participated in Experiment 2, and 17 in Experiment 3. They had not participated in other experiments concerned with contingency judgements. The apparatus and procedure for these experiments were identical to those described in Chapter 4.

4.8.1 Design

In each of the three experiments, participants were presented with three conditions with various relationships between unconditional and conditional ΔP. The event combinations for each condition are presented in Table 4.3. In Experiment 1, the three conditions had identical unconditional contingencies ($\Delta P_A = 0.5$ and $\Delta P_B = 0$), but were selected to produce an ascending pattern of $\Delta P_{A|B}$ and $\Delta P_{B|A}$ values and a descending pattern of $\Delta P_{A|\sim B}$ and $\Delta P_{B|\sim A}$. In Experiment 2, three conditions again had identical unconditional contingencies ($\Delta P_A = 0.5$ and $\Delta P_B = 0.72$), but were selected in order for $\Delta P_{A|B}$ to equal $\Delta P_{A|\sim B}$ and $\Delta P_{B|A}$ to equal $\Delta P_{B|\sim A}$. The conditional contingencies for both A and B resulted in an ascending pattern, while the unconditional contingencies for both A and B remained constant. Finally, in Experiment 3, the three conditions again had identical unconditional contingencies ($\Delta P_A = 0.5$ and $\Delta P_B = 0.72$), but were selected to produce an ascending pattern of $\Delta P_{A|B}$ and $\Delta P_{B|A}$ values and a descending pattern of $\Delta P_{A|\sim B}$ and $\Delta P_{B|\sim A}$ just as in Experiment 1, but with a stronger Cue B unconditional contingency in contrast to a non-contingent Cue B in Experiment 1. The Rescorla-Wagner model simulations were based on the same parameter values as indicated in Chapter 4, and based on 360 trials (rather than the 72 in Table 4.3) and 100 iterations. The nine conditions were simulated in turn and the resulting value for each cause after 360 trials is presented in Table 4.3.

4.9 Results

In order not to detract from the subject matter of this dissertation, the results from the three experiments are presented here simply for illustrative purposes to provide a brief sketch of how participants regard unequal conditional contingencies. Therefore, the summary statistics presented in Table 4.3 only report mean values after 72 trials. The noteworthy results include:

Table 4.3: Frequency of events in Experiment 1-3, mean ratings, actual and estimated (un)conditional ΔP values after 72 trials. The Rescorla-Wagner simulations were based on the same parameter values as indicated in Chapter 4, 360 trials, and 100 iterations.

	Experiment 1			Experiment 2			Experiment 3		
ABO	9	13	17	26	24	22	26	24	22
A∼BO	18	14	10	1	3	5	1	3	5
∼ABO	9	5	1	5	7	9	5	7	9
∼A∼BO	0	4	8	4	2	0	4	2	0
AB∼O	8	4	0	4	2	0	0	2	4
A∼B∼O	1	5	9	5	7	9	9	7	5
∼AB∼O	10	14	18	1	3	5	5	3	1
∼A∼B∼O	17	13	9	26	24	22	22	24	26
# of Trials	72	72	72	72	72	72	72	72	72
				Cue A					
ΔP_A	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
$\Delta P_{A\|B}$	0.06	0.44	0.95	0.03	0.22	0.36	0.5	0.22	-0.05
$\Delta P_{A\|\sim B}$	0.95	0.56	0.06	0.03	0.22	0.36	-0.05	0.22	0.5
$est\Delta P_A$	0.58	0.59	0.53	0.67	0.55	0.41	0.52	0.48	0.49
$est\Delta P_{A\|B}$	0.4	0.6	0.74	0.26	0.19	0.23	0.31	0.18	0.04
$est\Delta P_{A\|\sim B}$	0.76	0.59	0.31	0.32	0.33	0.32	0.17	0.28	0.38
Ratings	77.6	46	30.2	-25.3	8.2	0.5	-10.6	6.4	2.1
RW	0.5	0.52	0.49	0.02	0.22	0.36	0.23	0.21	0.21
				Cue B					
ΔP_B	0	0	0	0.72	0.72	0.72	0.72	0.72	0.72
$\Delta P_{B\|A}$	-0.42	-0.06	0.47	0.7	0.62	0.64	0.9	0.62	0.35
$\Delta P_{B\|\sim A}$	0.47	0.06	-0.42	0.7	0.62	0.64	0.35	0.62	0.9
$est\Delta P_B$	-0.02	-0.04	-0.08	0.77	0.76	0.67	0.73	0.67	0.73
$est\Delta P_{B\|A}$	-0.17	0	0.16	0.55	0.57	0.56	0.69	0.51	0.47
$est\Delta P_{B\|\sim A}$	0.19	-0.01	-0.27	0.61	0.71	0.65	0.55	0.62	0.81
Ratings	-11	-25.5	-50.7	50.1	55.4	48.1	42.6	19.2	66
RW	0.04	0.02	0.02	0.69	0.62	0.64	0.62	0.63	0.62

(1) the closer match between participants ratings and the contingencies conditional on the absence of the other cause ($\Delta P_{A|\sim B}$ and $\Delta P_{B|\sim A}$) rather than on the presence of the other cause ($\Delta P_{A|B}$ and $\Delta P_{B|A}$); (2) the corresponding data trends between the actual (un)conditional and estimated (un)conditional contingencies ($\Delta P \approx est\Delta P$); and (3) the intermediate positioning of the Rescorla-Wagner model simulation data, i.e., midway between the contingencies conditional on the absence of the alternative cause and the presence of the alternative cause.

The results from these three experiments suggest that participants' tend to rate the influence of each cause conditional on the absence of the other cause. This tendency is not reflected by the Rescorla-Wagner model. When the conditional contingencies are not equal, then the conditional ΔP account and the Rescorla-Wagner model make different predictions. In its current form, the conditional ΔP account does not specify the circumstances under which one conditionalises on the presence or absence of another cause. Cheng and Holyoak (1995) speculate that tests based on the absence of other causes are more informative than those in which other causes are present. Citing a personal communication with Patricia Cheng, Spellman et al. (2001) speculate that "...once subjects have reason to believe that both factors are potentially causal, they assess each cause only in the absence of the other cause" (p. 206). The results from the present experiments lend support to this conjecture. The Rescorla-Wagner model, on the other hand, seems to predict what is essentially the average of the conditional contingencies. However, an explanation for averaging conditional contingencies is not evident.

Chapter 5

Temporal Contiguity and Contingency Judgments: A Pavlovian Analogue

A watched pot never boils.

Irish Proverb

5.1 Preface

In October, 2002, Dr. Lorraine Allan was the first author on a paper presented at the 2002 meeting of the Pavlovian Society in Westwood, California. Following the conference, we were asked to submit an article to the journal of the Pavlovian Society: *Integrative Physiological and Behavioral Science.* Since the work presented at the conference was based on Tangen and Allan (2003), we decided to submit a paper based on our research on temporal contiguity. The chapter is reproduced from Allan et al. (in press) which was accepted for publication on March 25, 2003. Unlike the three previous chapters which measured cue-interaction between the two cues and a common outcome, this chapter only uses one cue and one outcome. Furthermore, rather than examining how knowledge of causal asymmetry affects causal judgements, in this chapter, we investigate how knowledge of temporal contiguity influences causal judgements. The results indicate that causal ratings were higher when the observer's expectations were congruent with the experienced delay than when the observer's expectations were incongruent with the experienced delay. While these data provide evidence for high-level processes, we discuss the predictions of the *temporal coding hypothesis*, an associative model which seems to account for this (seemingly) high level processing using a basic associative mechanism.

5.2 Abstract

Two experiments are reported that examine the role of temporal contiguity on judgments of contingency in a human analogue of the Pavlovian task. The data show that the effect of the actual delay on contingency judgment depends on the observer's expectation regarding the delay. For a fixed contingency between the cue and the outcome, ratings of the contingency are higher when the actual delay is congruent with the observer's expectation than when it is incongruent. We argue that our data can be understood within the context of the *temporal coding hypothesis.*

5.3 Introduction

There is considerable evidence of similarities between the operations that modulate the strength of conditioning in nonhuman animals and those that modulate the rating of the contingency between events by humans (see Allan, 1993). One of these similarities is the effect of temporal contiguity. It is well established in the animal literature that temporal contiguity is an important variable in both instrumental and Pavlovian conditioning (see Allan, Balsam, Church, & Terrace, 2002; Allan & Church, 2002). For example, increasing the delay between a response and reinforcement in an instrumental task decreases the rate of responding. Similarly, increasing the delay between a conditioned stimulus and an unconditioned stimulus in a Pavlovian task retards the acquisition of the conditioned response.

The studies that have examined the effect of temporal contiguity on ratings of contingency have used human analogues of the animal instrumental procedure (e.g., Buehner & May, 2002, in press; Reed, 1992, 1996; Shanks, 1989; Shanks & Dickinson, 1991; Shanks, Pearson, & Dickinson, 1989; Wasserman & Neunaber, 1986). In these instrumental studies, observers were required to perform an action, A, (e.g., tapping a key, pressing a button, pressing the space bar on a computer keyboard), and judge the extent to which the action was related to or caused the occurrence of an outcome, O, (e.g., illumination of a light, illumination of a triangle on a computer monitor, an explosion). Overall, these experiments found that the judged contingency between the action and the outcome decreased as the temporal delay between the action and the outcome was increased.

In an instrumental task, the observer is free to choose whether and when to respond. If an action occurs, then the outcome is presented with probability $P(O|A)$; if an action does not occur then the outcome is presented with probability $P(O| \sim A)$. The contingency between the action and the outcome, ΔP, is defined as the difference between these two independent conditional probabilities:

$$\Delta P = P(O|A) - P(O| \sim A) \tag{5.1}$$

Temporal contiguity is varied by inserting a delay between the action and the outcome. This design has a number of difficulties, however. The overall probability of an action, $P(A)$, is determined by the observer, rather than by the experimenter, and has been shown to vary among observers and more importantly between delay conditions. Generally, $P(A)$ decreased as the temporal delay between the action and the outcome was increased (see Buehner &

Table 5.1: Standard 2 × 2 contingency matrix for the human analogue of the Pavlovian task

		Outcome		
		O	∼O	
Cue	C	a	b	a+b
	∼C	c	d	c+d
		a+c	b+d	

May, in press). Also, although the observer is allowed to respond during the delay, in many of the human studies these extra responses did not result in an outcome. As Buehner and May showed, these extra responses during the delay can result in the actual values of $P(O|A)$ and $P(O| \sim A)$ differing from the intended values. Thus, actual ΔP can change as a function of delay, and the variation in rating with delay might reflect a change in ΔP rather than a change in temporal contiguity. While control groups were often included in the design of the experiments using the instrumental analogue, an alternative approach would be to use a task where such confounds are removed. This could be accomplished by varying temporal contiguity in a human analogue of the Pavlovian task.

Although many of the early studies of human contingency judgments used the human analogue of the instrumental task, later studies did switch to the human analogue of the Pavlovian task (see Dickinson, 2001). Surprisingly however, the effect of temporal contiguity on contingency judgments has not been investigated in a Pavlovian situation where the cue-outcome combinations are experience by the observer. While Hagmayer and Waldmann (2003) were interested in the influence of temporal assumptions about cue-outcome relationships, they used a described format rather than an experienced format. In their experiments, the cue-outcome combinations on each trial were listed on a sheet of paper. In this type of presentation, the actual delay between the cue and the outcome can only be described and is not experienced in real time. In the experiments reported in the present paper, we vary temporal contiguity in real time in the human analogue of the Pavlovian task.

Table 5.1 presents the standard 2 × 2 contingency matrix for the human analogue of the Pavlovian task. In such tasks, the cue is either present (C) or absent (∼C) and the outcome is either present (O) or absent (∼O). The four cells (a, b, c, d) represent the joint frequency of occurrence of the four possible cue-outcome combinations. The contingency between the cue and the outcome, ΔP, is the difference between two independent conditional probabilities (Allan, 1980). Referring to Table 5.1,

$$\Delta P = P(O|A) - P(O| \sim A \sim B) = \frac{a}{a+b} - \frac{c}{c+d} \tag{5.2}$$

The effect of temporal contiguity could be investigated by manipulating the cue-outcome interval or delay.

Table 5.2: Cell frequencies and conditional probabilities for Experiments 1 and 2

	Experiment 1				Experiment 2	
CO	3	7	13	17	12	8
C∼O	17	13	7	3	4	8
∼CO	17	13	7	3	4	8
∼C∼O	3	7	13	17	12	8
# of Trials	40	40	40	40	32	32
$P(O\|C)$	0.15	0.35	0.65	0.85	0.75	0.5
$P(O\|\sim C)$	0.85	0.65	0.35	0.15	0.25	0.5
ΔP	-0.7	-0.3	0.3	0.7	0.75	0

5.4 Experiment 1

5.4.1 Method

Observers

The observers were 30 undergraduate students enrolled in Psychology courses at McMaster University who participated for course credit. They had not participated in other experiments concerned with contingency judgments. An equal number (n = 15) were randomly assigned to each of two delay groups (0.4 and 2 sec).

Apparatus

Observers performed the experiment on Power Macintosh computers located in separate rooms. The experiment was programmed in MetaCard 2.3.1. The stimuli were identical to those used in Tangen and Allan (2003). The cue was a chemical and the outcome was a bacterial strain. The presence of a cue was indicated by a colored three-dimensional animation of a chemical spinning on its axis and the absence of a cue was indicated by a faded unmoving grayscale picture of the chemical. Similarly, the presence of the outcome was indicated by a colored animation of moving bacteria and the absence of the outcome was indicated by a faded grayscale picture of stationary bacteria. There were five chemicals and five bacterial strains. Each of the chemicals and strains were randomly assigned fictitious names from a set of five chemicals and five bacteria. The colored moving images (indicating the presence of cues and outcomes) were accompanied by a name (e.g., Chorbine Present). In contrast, the faded unmoving images (indicating the absence of cues and outcomes) were not accompanied by text.

Procedure

The instructions for the experiment were presented to the observer on the computer monitor (see Appendix D for the full instructions). In brief, the observer was told that scientists have recently discovered four strains of bacteria that exist in the mammalian digestive system. For each strain, the scientists were testing whether a chemical aids in, interferes with, or has no effect on a strain's survival. To do this, a strain of bacteria was first placed in culture (petri dishes). After that, a chemical might be added to the bacterial culture. The scientists then verified whether or not the bacterial sample survived. The observer was also shown the rating scale that they would use to rate the effectiveness of the chemical on the survival of the bacteria. After reading the instructions, the observer was shown a summary screen of the four cue-outcome combinations. Eight practice trials were then presented where each of the four cue-outcome combinations was presented twice in random order.

The two cue-outcome delays (0.4 sec and 2.0 sec) were varied between subjects, and the four values of ΔP $(-.7, -.3, .3, .7)$ were varied within subjects. The frequencies and conditional probabilities for each ΔP are shown in Table 5.2. An experimental session consisted of four blocks of 40 trials each. The order of the four ΔP values was randomly determined for each observer over the four blocks.

The observer initiated a block of trials by clicking on the Begin button on the computer screen. A trial consisted of one of the four cue-outcome combinations: chemical added and bacteria survived (CO), chemical added and bacteria did not survive (C∼O), chemical not added and bacteria survived (∼CO), and chemical not added and bacteria did not survive (∼C∼O). The next trial was initiated by a mouse click on the Next Trial button.

During each block of trials, the observer rated how strongly the chemical affected the survival of the bacterial strain after trial 20 and again after trial 40. The ratings were made on a horizontal scrollbar that ranged from -100 (chemical has a very strong negative effect on the bacteria's survival) on the left to +100 (chemical has a strong positive effect on the bacteria's survival) on the right, and was anchored at 0 in the middle. Observers made their ratings by moving a horizontal scrollbar left and right with the mouse.

Each block of 40 trials was clearly labeled as separate scientific experiments with different chemical and bacteria images and names. For each observer, one chemical and one bacterial strain was randomly assigned to each of the four ΔP values. The remaining chemical and the remaining bacterial strain were used in the practice trials.

5.4.2 Results and Discussion

Figure 5.1 shows the mean ratings as a function of ΔP for each delay. The ratings at trial 20 are seen in Figure 5.1a and the ratings at trial 40 are seen in Figure 5.1b. For both delays and at both trials, ratings are an orderly function of ΔP, being negative for the negative contingencies and positive for the positive contingencies. At neither trial, do the ratings appear to differ for the two delays. A 2 × 2 × 4 mixed-design ANOVA was conducted on the ratings with delay (0.4 and 2.0) as a between-subject variable, and trial (20 and 40) and contingency $(-.7, -.3, .3, .7)$ as within-subject variables. The only main effect that was significant was contingency, $F(3, 84) = 72.86$, $p < .001$. The only other significant result was

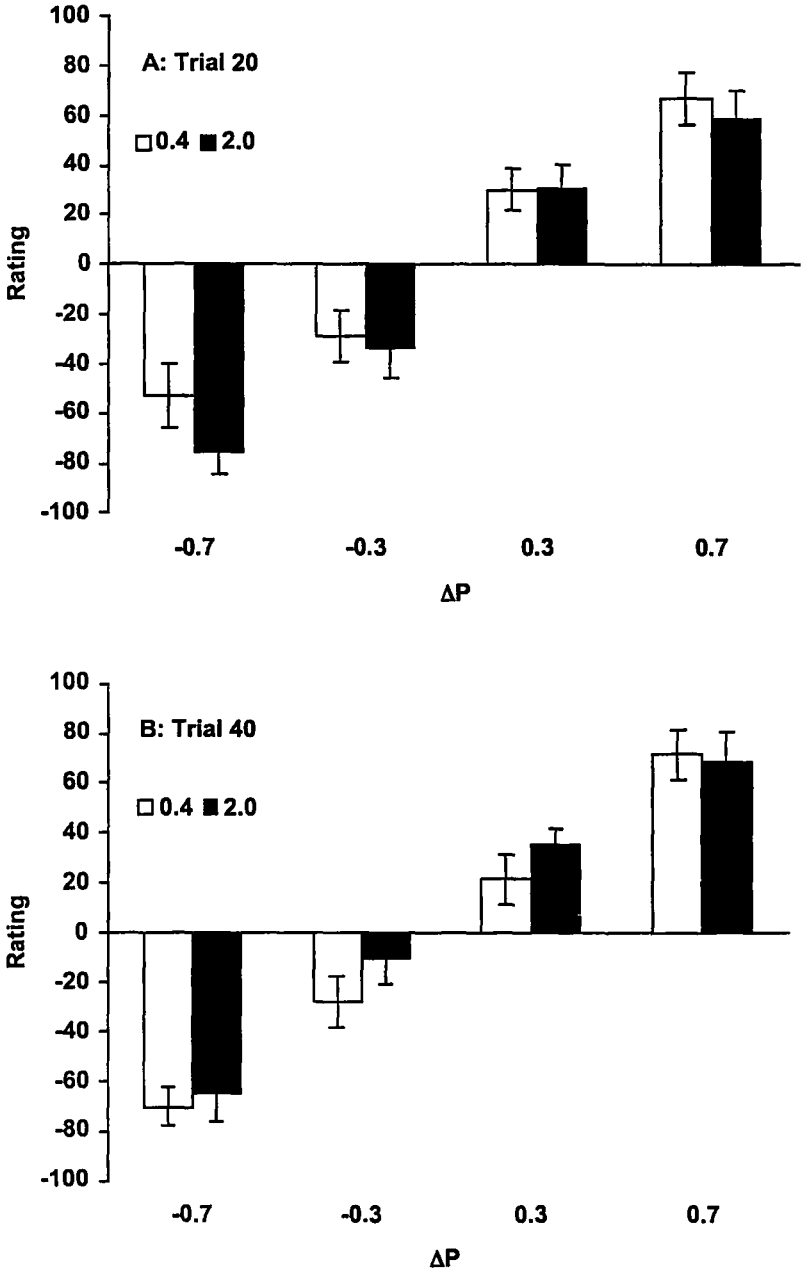

Figure 5.1: Mean ratings in Experiment 1 as a function of ΔP for each delay (unfilled bars for .4 sec and filled bars for 2.0 sec). The ratings at trial 20 are seen in Figure 5.1a and the ratings for trial 40 are seen in Figure 5.1b.

the interaction between trial and delay, $F(1, 28) = 4.63$, $p < .05$. At trial 20, the ratings were higher for the 0.4 sec delay than for the 2.0 sec delay, whereas at trial 40 the reverse was the case. The Tukey test indicated, however, that at neither trial were the ratings significantly different for the two delays, $ps < .05$.

Our manipulation of delay in Experiment 1 was without effect. Recent research reported by Buehner and May (2002, in press) suggests a plausible explanation for our null result. They showed that the effect of delay in the instrumental task depended on the cover story describing the action and the outcome. If the cover story indicated that the action produced an immediate outcome, then ratings were higher for a short delay than for a long delay. However, if the cover story indicated that the action produced a delayed outcome, then ratings did not differ as a function of delay. Such data suggest that the observer's expectation about the delay might be an important variable.

Our cover story in Experiment 1 was equally compatible with both the short and the long delay. There was no information in the cover story that would lead the subject to think that the effect of the chemical on the survival of the bacteria should occur immediately or should be delayed. Thus it is plausible that the observers who experienced the .4 sec delay interpreted the cover story to expect a short delay and the observers who experienced the 2.0 sec delay interpreted the cover story to expect a long delay. In Experiment 2, we look at the effect of cover story in the Pavlovian task. We use two different cover stories, one designed to induce an immediate expectation of the outcome and the other to induce a delayed expectation of the outcome. We cross the delay described in the two cover stories with the actual delay between the cue and the outcome presented to the observer.

5.5 Experiment 2

As we noted earlier, Buehner and May (2002, in press) showed that the effect of delay in the instrumental task depended on the cover story describing the action and the outcome. In Experiment 2, we adapted the cover stories used by Buehner and May (2002, Experiment 3) to the Pavlovian task. In the Buehner and May experiment, the observer's task was to determine whether triggering a "FIRE" button produced an explosion in a training range. In the immediate cover story, the observer was told that the FIRE button was a remote control detonator which, when fired, set off a mine in the training range immediately upon bring fired. In the delay cover story, the observer was told that the FIRE button was a grenade launcher which, when fired, sent shells into the training range. Since these shells had to travel, there would be a delay between pressing the FIRE button and the resulting explosion. In addition to varying the cover story (immediate and delay), Buehner and May varied the actual delay between the pressing of the FIRE button and the explosion (0 sec and 5 sec). They found a significant interaction between cover story and delay. At the 5 sec delay, ratings were significantly higher in the delay cover story than in the immediate cover story. That is, when there was an actual delay between the action and the outcome, ratings were higher when observers expected a delay than when they did not.

In the Buehner and May (2002) instrumental task, the observer decided whether and when

to press the FIRE button. In our experiment, we used discrete trials and whether the FIRE button was pressed on a trial was preprogrammed. We varied the actual delay between the cue and the outcome (0 sec and 5 sec) presented. We examined the interaction between cover story and cue-outcome delay for two values of ΔP (0 and 0.5).

5.5.1 Method

Observers

The observers were 52 undergraduate students enrolled in Introductory Psychology at McMaster University who received course credit and who had not participated in other experiments concerned with contingency judgments. An equal number (n = 13) were randomly assigned to each of four groups.

Apparatus

The apparatus was the same as in Experiment 1, and the experiment was again programmed in MetaCard 2.3.1. There were four movie clips corresponding to the four trial types presented in Table 5.1. For each of the four clips, a computer rendered animation (created in Poser 4.0.1) of a military officer is presented. The officer either presses the FIRE button (C) or not (~C), resulting in an explosion on the horizon (O) or not (~O).

Procedure

The instructions for the experiment were presented to the observer on the computer monitor (see Appendix D for the full instructions). There were two instruction sets, one for the immediate cover (IC) story and one for the delay cover (DC) story. In the IC story, the FIRE button was described in the context of a remote-control detonator and the observer was told that the mine explosion was immediate. In the DC story, the FIRE button was described in the context of a grenade launcher and the observer was told that there was a delay of a few seconds between the pressing of the FIRE button and the resulting mine explosion. Both sets of instructions explained that the device was still in the experimental phase, and therefore pressing the FIRE button did not always result in an explosion and also that an explosion might occur even when the button was not pressed. The instructions also described the scale that they would use to rate the relationship between the FIRE button being pressed and the mine exploding. After reading the instructions, the observer was shown a summary screen of the four cue-outcome combinations.

Half the observers received the IC story and half received the DC story. Two cue-outcome delays were used, 0 sec and 5 sec. Each observer experienced both delays. In each cover-story group, half the observers experienced the 0 sec delay first, followed by the 5 sec delay (0/5 order). The order was reversed for the remaining observers (5/0 order). There were two values of ΔP, 0 and 0.5. The frequencies and conditional probabilities for each ΔP value are shown in Table 5.2. The order of the two ΔP values at each delay was randomly determined. In

summary, there were two between variables (cover story and order) with two levels each, and two within variables (delay and ΔP) with two levels each.

An experimental session began with eight practice trials where each of the four cue-outcome combinations was presented twice in random order. On these practice trials, $\Delta P = .5$, and the delay was the same as that programmed for the first delay to be experienced. The practice trials were followed by four experimental blocks of 32 trials each. The observer initiated a block of trials by clicking on the "Begin" button on the computer screen. A trial consisted of one of the four cue-outcome movie clips: button pressed and explosion, button pressed and no explosion, button not pressed and explosion, button not pressed and no explosion. The next trial was initiated by a mouse click on the "Next Trial" button. At the end of each block, the observer rated the relationship between pressing the FIRE button and the explosion of a mine. The ratings were made on a horizontal scrollbar that ranged from 0 (the FIRE button had no effect on causing the explosion) to +100 (the FIRE button was a perfect cause of the explosion). Observers made their ratings by moving a horizontal scrollbar left and right with the mouse.

5.5.2 Results and Discussion

Figure 5.2 shows mean ratings as a function of delay (0 and 5 sec) for the two cover stories (IC and DC) at each of the two ΔP values (0 and 0.5). Figure 5.2a presents the data for the 0/5 order and Figure 5.2b presents the data for the 5/0 order. Both figures indicate that for each ΔP value, there is an interaction between cover story and delay. Ratings are higher when the cover story and the delay were congruent (IC with 0 sec delay and DC with 5 sec delay) than when the cover story and the delay were incongruent (IC with 5 sec delay and DC with 0 sec delay). A $2 \times 2 \times 2 \times 2$ mixed-design ANOVA was conducted on the ratings, with cover story (IC and DC) and order (0/5 and 5/0) as between-subject variables, and delay (0 sec and 5 sec) and ΔP (0 and 0.5) as within-subject variables. The main effect of ΔP was significant, $F(1, 48) = 120.08$, $p < .001$. There was a significant interaction between cover story and delay, $F(1, 48) = 34.59$, $p < .001$. The Tukey test indicated that at the 0 sec delay ratings for IC (55.65) were higher than for DC (42.46), $p < .01$ and that at the 5 sec delay ratings for IC (36.73) were lower than for DC (54.87), $p < .001$. The Tukey test also indicated that for IC, the ratings were higher at the 0 sec delay (55.65) than at the 5 sec delay (36.73), $p < .001$ and that for DC, the ratings were lower at the 0 sec delay (42.46) than at the 5 sec delay (54.87), $p < .01$. The Tukey test also confirmed that the two congruent combinations (IC with 0 sec and DC with 5 sec) did not differ, $p > .05$, and that the two incongruent combinations (IC with 5 sec and DC with 0 sec) did not differ, $p > .05$.

The ANOVA also revealed a significant three-way interaction between cover story, delay, and ΔP, $F(1, 48) = 13.57$, $p < .001$. The interaction of cover story and delay was more pronounced for $\Delta P = 0.5$ than for $\Delta P = 0$. The three-way interaction between cover story, delay, and order was also significant, $F(1, 48) = 4.37$, $p < .05$. The interaction of cover story and delay was more pronounced for the 0/5 order than for the 5/0 order.

Our data extend the findings of Buehner and May (2002), which showed an interaction of cover story and delay on contingency ratings in the instrumental task, to the Pavlovian task.

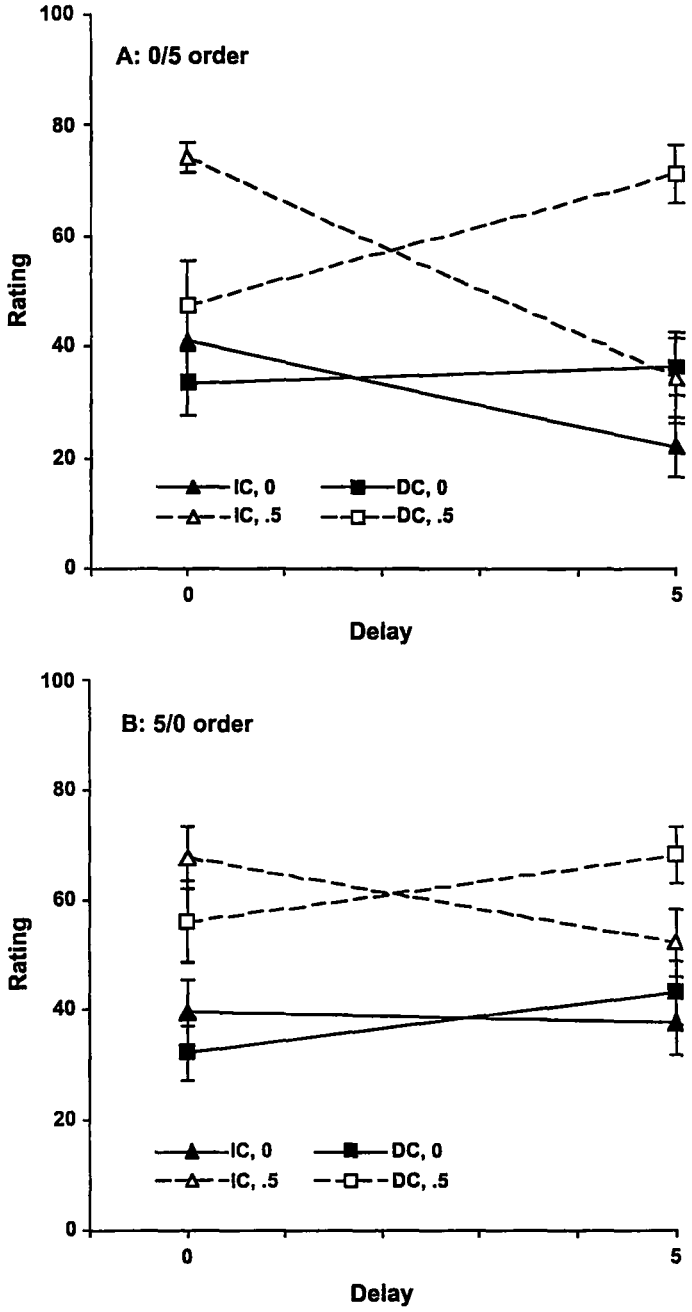

Figure 5.2: Mean ratings in Experiment 2 as a function of delay for the two cover stories (triangles for IC and squares for DC) at each ΔP value (filled symbols for $\Delta P = 0$ and unfilled symbols for $\Delta P = .5$). The ratings for the 0/5 order are in Figure 5.2a and the ratings for the 5/0 order are in Figure 5.2b.

Buehner and May demonstrated an effect of cover story at 5 sec but not at 0 sec. Specifically, ratings were higher for the DC cover story than for the IC cover story at the 5 sec delay, but did not differ at the 0 sec delay. Our data provide even stronger evidence for the role of cover story. We find an effect of cover story at both delays, and the direction of the effect is different at the two delays. It should be emphasized that the stimulus events were identical for the two cover stories and that the instructions were similar with only a few crucial words being changed. We also show that the interaction between cover story and delay occurs not only when there is a relationship between the cue and the outcome ($\Delta P = 0.5$) but also when there is no relationship ($\Delta P = 0$).

5.5.3 Discussion

The data from Experiment 1 suggested that increasing the delay between the cue and the outcome is ineffective if the delay is consistent with the observer's expectations. The data from Experiment 2 confirmed that this was the case. These data demonstrated that when the delay is congruent with the observer's expectation, ratings at a 5 sec delay do not differ from ratings at 0 sec delay. The data also indicate that ratings are higher when the cover story and the delay are congruent (IC with 0 sec delay and DC with 5 sec delay) than when the cover story and the delay are incongruent (IC with 5 sec delay and DC with 0 sec delay).

Buehner and May (2002, in press) argued that an interaction between cover story and delay is not consistent with associative accounts of contingency judgments. Associative models postulate that judgments are determined by associative links or connections which are formed between contiguously-presented cues and outcomes. The associative model most frequently used to account for human contingency judgments is the Rescorla-Wagner model (Rescorla & Wagner, 1972), although recently the Pearce generalization model (Pearce, 1987) has gained popularity. In these associative models, temporal contiguity contributes to the strength of the association but does not become part of that association. That is, temporal factors serve only a facilitative role in the formation of associations. The closer the cue and the outcome are in time during training, the more robust the resulting association is assumed to be. The organism, however, acquires no representational knowledge about the temporal relationship between the cue and the outcome.

While temporal factors serve only a facilitative role in the formation of associations in some associative models (e.g., Rescorla & Wagner, 1972), this is not the case for all associative models (e.g., Wagner, 1981; Miller & Barnet, 1993). In fact, a prime motivating factor for the development of more recent associative models was to directly encompass the role of temporal factors in the learning process. For example, the temporal coding hypothesis (e.g., Miller & Barnet, 1993) questioned the assumption that organisms do not learn about the temporal properties of the stimulus events. According to the temporal coding hypothesis, the temporal relationship between the cue and the outcome is automatically encoded as part of the content of an association and plays a critical role in determining the response. Recently, Savastano and Miller (1998) provided an overview of the accumulating evidence indicating that organisms do acquire temporal information in a wide variety of Pavlovian paradigms. Organisms can superimpose temporal maps when elements common to these maps are presented together,

even when the elements were trained separately. That is, temporal information from different training situations can be integrated.

The evidence summarized by Savastano and Miller (1998) in support of the temporal coding hypothesis stems mainly from the animal learning literature. An animal learning experiment analogous to our Experiment 2 would cross delay in Phase 1 (0 sec or 5 sec) with delay in Phase 2 (0 sec or 5 sec). In their review, Savastano and Miller do not report animal data from this design, but they do discuss similar designs. For example, Savastano, Yin, Barnet, and Miller (1998) described an experiment modeled on the Hall and Pearce (1979) CS-preexposure effect. In the usual CS-preexposure experiment, the CS is presented alone (i.e., without the US) first. Following this preexposure phase, the CS and the US are paired. CS-preexposure weakens the strength of the CR compared to appropriate control groups. In the preexposure phase of the Hall and Pearce situation, the CS was paired with a low intensity version of the US (CS $\rightarrow$ US_{weak}). This was followed with the same CS being paired with a high intensity version of the US (CS $\rightarrow$ US_{strong}). Hall and Pearce found that preexposure to the weak version of the US also resulted in an attenuated CR. Using the Hall and Pearce preexposure design, Savastano et al. (1998) varied the delay between the CS and the US in both phases. They showed that the size of the CS-preexposure effect did not depend on the absolute values of the CS-US delays, but rather on whether the temporal relationship between the CS and the US in the two phases was congruent or incongruent.

The prediction of the temporal coding hypothesis for an experiment that crossed Phase 1 delay (0 sec or 5 sec) with Phase 2 delay (0 sec or 5 sec) is clear – the conditioned response should be stronger when the cue in the two phases shares the same temporal relationship with the outcome (the two congruent conditions) than when they share different temporal relationships (the two noncongruent conditions). The cover story in our Experiment 2 would be analogous to Phase 1 in the animal learning experiment. Rather than actually experiencing the delay, the observer "learns" the delay by reading the cover story (IC or DC). This initial learning phase is followed by a second phase where the observer actually experiences a delay that is congruent or incongruent with the temporal relationship described in the cover story. The temporal coding hypothesis, applied to our Experiment 2, would predict that, because ratings are influenced by temporal relationships, the ratings should be higher in the two congruent conditions than in the two noncongruent conditions. Thus the temporal coding hypothesis, which falls into the category of associative models, can encompass our results.

Our experiments were not designed as a test of the temporal coding hypothesis. Rather, their purpose was to examine the role of temporal contiguity in the human analogue of the Pavlovian task. What we wish to emphasize, however, is that the interaction between cover story and delay that we observed, and that was also reported by Buehner and May (2002, in press), is compatible with an associative account of human contingency judgments. Researchers who argue against an associative account of contingency judgments tend to base their evaluation on a particular associative account and then generalize to all associative accounts (see Allan, 2003). Buehner and May (in press) argued that "...assumptions [of associative models] are too rigid in that they clearly state that delays should always hinder performance; instructional effects of the kind we found in our experiments cannot be accounted for under such a theory". This is not the case.

Chapter 6

General Discussion

> The belief that we can start with pure observations alone, without anything in the nature of a theory, is absurd.
>
> *Popper (1972)*

The objective of this dissertation has been to assess the circumstances in which our general knowledge of causal direction and temporal delay guides judgements of causality. The results reported in Chapters 2 and 3 suggest that knowledge of causal direction modulates judgements depending on the nature of task, and the results reported in Chapter 5 suggest that knowledge of temporal delay modulates causal judgements in each of the conditions tested. The results from each chapter will be discussed in turn.

In Chapter 2, Experiments 1 and 2 demonstrate that knowledge of causal structure interacts with the level of cue-interaction by uncorrelating causal order and the number of cues and outcomes that have typically been confounded in the literature. The results from Experiment 3 suggest that causal knowledge only influences cue-interaction in causal ratings, but not trial-by-trial prediction responses. Lastly, in Experiment 4, the causal model effect dissipated as the trials progressed. We propose in Chapter 2 that both low-level (associative) factors and high-level (causal model) factors influence causal assessment depending on what is being asked about the events, and participants' experience with those events. We provided several predictions as to the circumstances under which we would expect each of the two factors to be more heavily weighted.

The three experiments in Chapter 3 demonstrate that, in general, knowledge of causal direction influences causal ratings, but not trial-by-trial predictions. We note that while the causal description of the events influence ratings when comparing the common-effect (2C-1E) and common-cause (2E-1C) scenarios, the results do not exactly track the predictions made by causal-model theory. It is evident that causal knowledge had an influence, but not to the extent assumed by causal-model theory. The influence of the causal model was particularly evident in Experiment 3 in which participants were asked to provide an integrative causal rating. We derived this manipulation from one of our predictions in the previous chapter. Experiments 1 and 2, however, examined the clarity of the causal model and the wording

of the test question respectively, resulting in less sensitivity to the causal model. In each experiment, sensitivity to the causal model was not evident in the trial-by-trial prediction responses.

Overall, when one examines the data from the seven experiments presented in Chapters 2 and 3, it is evident that participants are sensitive to the causal description of the events. According to an associative model, when two cues precede a common outcome, the cues will compete to be associated with the outcome *despite* their causal description. Conversely, when one cue precedes two outcomes, the outcomes will not compete. Therefore, an associative account would predict an identical pattern of results for the 2C-1E and 2E-1C scenarios, and an identical pattern of results for the 1C-2E and 1E-2C scenarios. The results from Experiments 1 and 2 in Chapter 2 indicate that participants' do not judge them to be the same. Their ratings interact with the direction of the causal relationship. Ratings for 2E-1C tend to be higher than for 2C-1E, and 1E-2C ratings tend to be lower than for 1C-2E. Furthermore, in Experiments 3 and 4 of Chapter 2, and Experiments 1-3 of Chapter 3, even though assessments of the 2C-1E and 2E-1C scenarios are more similar, judgements of 2E-1C are consistently higher than for 2C-1E. An associative account alone cannot account for this difference.

On the other hand, when one examines the data from the seven experiments presented in Chapters 2 and 3, participants also seem to be sensitive to the associative nature of the events. According to causal-model theory, when two causes produce a common effect, the causes will interact, and when two effects result from a common cause, the effects will not interact. Causes will interact and effects will not *despite* the order that they are presented. Therefore, causal-model theory would predict an identical pattern of results for the 2C-1E and 1E-2C scenarios, and an identical pattern of results for the 2E-1C and 1C-2E scenarios. The results from Experiments 1 and 2 in Chapter 2 indicate that while participants' rate them similarly, their ratings for 1C-2E tend to be higher than for 2E-1C, and 2C-1E ratings tend to be lower than for 1E-2C. Furthermore, in Experiments 3 and 4 of Chapter 2, and Experiments 1-3 of Chapter 3, assessments of the 2C-1E and 2E-1C scenarios tend to be very similar. A causal-model account alone cannot account for this difference.

The results from Chapters 2 and 3 suggest that our causal assessments are best explained by the joint contribution of high and low-level processing whereby causal judgements are influenced both by causal knowledge and by the associative nature of the events. The debate in the literature implies that causal judgements are *either* concept-driven or data-driven, but not both. For example, Cobos et al. (2002) conclude that "... the circumstances under which a [causal-model theory]-like influence is observed are quite restricted and, correspondingly, that the circumstances under which causal inferences may be mediated by the operation of an associative learning mechanism are quite broad." (p. 344). In contrast, Waldmann and Holyoak (1992) conclude by inquiring whether "lower-order associative learning should be reduced to higher order causal induction" (p. 235). Given the results from Chapters 2 and 3, it may more productive to consider causal assessment as driven by the interaction between data-driven (associative) and concept-driven (causal-model) processes. It can be argued that positing a completely associative or causal knowledge-based processing model will fail to capture the complex interactive nature of causal judgement. Alternatively, an interactive

approach to causal judgement appears to offer a better explanation of the cognitive processes which seems to be at work here.

In Chapters 2 and 3, participants rated a moderately positive cause as less predictive when it was paired with a strong predictor than to when it was paired with a weak predictor. Specifically, when two causes preceded a common effect (2C-1E), and the contingency of Cause A with the common effect was moderately contingent ($\Delta P_A = 0.5$), then causal ratings of A decreased as the contingency of Cause B with the common effect increased (e.g., ΔP_B = 0, 0.25, 0.5, 0.75, 1). However, the strength of contingency referred to throughout this thesis (e.g., "strong", "moderate", and "weak" contingencies) and in the literature refer to the unconditional contingencies. In Chapters 2 and 3, we indicated that "several different frequencies can be selected to fill the eight cells of the 4×2 matrix each resulting in various combinations of unconditional and conditional ΔP values. The frequencies for Experiment 2 [Chapter 2] (shown in Table 2.2) were selected to produce a descending pattern of conditional ΔP_A values while maintaining identical unconditional ΔP_A values. As well, they were selected so the unconditional and conditional ΔP_B values would be as closely matched as possible. The ΔP_B values were therefore selected only for their influence on the conditional ΔP_A values. The frequencies were also selected so the respective conditional contingencies for A and B would be identical where $\Delta P_{A|B} = \Delta P_{A|\sim B}$ and $\Delta P_{B|A} = \Delta P_{B|\sim A}$ resulting in the symmetry observed in the two columns of the 4×2 contingency matrix for each of the five conditions (see Spellman, 1996b, Property 4)."

In Chapter 4, we demonstrate that when two causes result in one effect, the relative "strength of contingency" in cue-interaction effects, is based on conditional rather than unconditional contingencies. That is, participants' ratings tracked the conditional rather than uncondtional contingencies as predicted by the conditional ΔP account (Spellman, 1996b). We also demonstrated that the Rescorla-Wagner model makes the same predictions as conditional ΔP at asymptote given Properties 1 and 4 made by Spellman (1996b).

Again, while the two models make identical predictions about cue-interaction given the assumptions above, they make very different predictions about how these judgements are formed. In particular, the Rescorla-Wagner model updates the associative strength accrued to each cue on each trial. Therefore, the model makes specific predictions about the acquisition function for each cue that varies depending on the ordering of the events. When researchers simulate the event frequencies for certain experimental conditions, they generally input random event sequences for a number of iterations. What they are essentially doing is collapsing over the order that the events were presented thereby disregarding the individual trial sequences. Under these circumstances (and given the assumptions above), the Rescorla-Wagner model and conditional ΔP make identical predictions. In computing conditional ΔP, one simply "counts up" the number of each event type and computes each of the probabilities. Basing these statistics on frequency information disregards the order that those frequencies were obtained. Furthermore, the Rescorla-Wagner model allows one to make pre-asymptotic predictions while conditional ΔP does not. For example, ΔP is computed identically whether it is based on 10 trials or 1000 trials. Because the Rescorla-Wagner model divvies up λ (the sum of associative strength) among the events over a number of trials, the associative strength accumulated pre-asymptotically can vary greatly from the associative strength accumulated

at asymptote. Therefore, the Rescorla-Wagner model predicts trial effects, while conditional ΔP does not. As discussed in Chapters 2 and 4, participants' judgements varied across trials which cannot be accounted for by conditional ΔP.

In Chapter 5 we examined how assumptions of temporal delay influenced causal assessments. Specifically, in Experiment 1, the same chemical and bacteria materials were used as in each of the other experiments. In this case, however, only one chemical affected a single strain of bacteria. The independent measure for this experiment was the delay between the offset of the chemical and the onset of the bacteria. Previous experiments that have used an operant procedure (in which participants initiate the cause), and animal learning experiments that have used a similar Pavlovian design, suggest that participants' causal assessments should decrease as the temporal delay between the events increase. However, the temporal delay in this case had no influence on participants' causal ratings. There was no significant difference between judgements of events with a short delay and events with a long delay. We argued that the materials we used were consistent with *both* a short and a long delay. Therefore, participants' expectations were congruent with the events presented to them, resulting in equivalent ratings in each delay condition.

To further investigate the interaction between temporal knowledge and causal judgements, we designed a second experiment with a new set set of materials. In this case, participants were asked to rate how effective a button press was in producing an explosion some distance away. Given one of two cover stories, one group of participants expected the temporal delay between the events to be short, and another group expected the delay to be long. Participants were then presented with a series of trials that were either congruent with their expectations, or incongruent with their expectations. Causal ratings were higher if the observed events were congruent with their expectations then if the events were incongruent with their expectations. This interaction between expectation and observation was evident for both moderately contingent ($\Delta P = 0.5$) and non-contingent ($\Delta P = 0$) relationships.

According to the Rescorla-Wagner model, temporal contiguity serves only a facilitative role in the formation of associations. The closer events occur in time, the more robust the resulting association is assumed to be. However, the temporal coding hypothesis is unlike most associative models in that it encodes the temporal (in)congruity between different training situations. If the temporal interval of a learning phase is congruent with the temporal interval of another, then the conditioned response should be stronger than if the temporal intervals are incongruent between the learning phases. The role of temporal congruity according to the temporal coding hypothesis shares common utility with the role surmised by Hagmayer and Waldmann (2003). However, Hagmayer and Waldmann argue that knowledge of temporal intervals serves an identical role as knowledge of causal directionality, guiding the selection of statistical indicators of causal strength. Specifically, Hagmayer and Waldmann (2003) provide evidence that temporal assumptions influence the selection of relevant events that enter contingency estimates. In their second experiment, participants tended to base their contingency estimates only on events that were relevant to the particular temporal assumption. Events that occurred too long or too short before the effect were neglected as candidate causes. If they had no assumptions about temporal intervals, then all the events were considered. The results from Chapter 5 imply that if participants' expectations of the temporal delay were

incongruent with the actual delay, then they would likely disregard the button as a candidate cause of the explosion. Therefore, on trials in which the explosion occurred in isolation (∼AO) or when the button was pressed without an explosion (A∼O) they may have attributed the causal influence to whatever else was causing the explosion, thereby over weighting cells b and c in Figure 4.1 of Chapter 4, and under weighting cells a and d. In the congruent condition, the opposite would be true. However, in Experiment 1, our participants had no assumptions of temporal intervals, then perhaps they considered each event type equally by weighting cells a-d evenly.

Hagmayer and Waldmann (2003) assert that knowledge of temporal intervals influence causal judgements in much the same way as knowledge of causal asymmetry, in that they guide the choice of appropriate statistical indicators for causality. Perhaps the constraint of the space of possible relationships is a simplification strategy (Jennings, Amabile, & Ross, 1982b; Nisbett & Ross, 1980; Gigerenzer & Todd, 1999). For example, if we return to Ruth and Arthur introduced in Chapter 1, Ruth was faced with number of viruses and asked to predict the onset of a disease. Arthur was faced with a number of symptoms and asked to diagnose a particular illness. Given the data presented in Chapters 2 and 3, we know the fact that "Ruth is predicting" and "Arthur is diagnosing" influences how they regard each virus/symptom. Not only does causal asymmetry affect their judgements, but knowledge of causal direction significantly reduces the complexity of the learning situation. If Arthur's patient complains of a number of symptoms, he is justified in supposing that they originated from a common cause. In contrast, Ruth is justified in testing each virus separately to determine the likelihood of each resulting in the disease.

As Reichenbach (1956) pointed out, if both lamps in a room go out suddenly, or if several actors in a stage play fall ill showing symptoms of food poisoning, then we look for a common cause. It is possible that the bulbs burned out simultaneously, or that the actors fell ill independently for different reasons, but it is more plausible to suggest a common cause. Our knowledge of causal direction allows us to explain coincidences by appealing to a common cause, and conditionalise by appealing to a common effect. These examples are simplistic, but confronted with a web of causally related events (e.g., epidemiology, economics, and experimentation), reducing the complexity of the event space would clearly be adaptive.

Perhaps temporal expectation also serves as a simplification strategy. If Arthur knows that the spoiled fish his patient ate has an incubation period of two days, then he can discount the reported symptoms long before and long after the two day period as resulting from an unrelated cause. Faced with a large number of related events, Arthur's temporal assumptions constrain the event space to those that occur within a particular timeframe.

Throughout this dissertation, I have examined the role of expectation on judgements of causality. In Chapter 5, we investigated the congruity between participants' temporal expectations and observations. They observed events that were either congruent or incongruent with their temporal assumptions. However, the issue of congruity in expectations of causal direction has not been discussed. Expectation of causal asymmetry seems to be of a different kind than expectation of temporal delay. Knowledge of causal asymmetry rests solely on the assumption that causes interact and effects do not. Identifying common cause and common effect causal structures will significantly reduce the complexity of the learning task.

However, these causal structures are dichotomous and, therefore, do not lend themselves well to the notion of congruity. The observed events have one causal structure or another which either confirm one's expectations or not. Temporal relationships, on the other hand, can be either qualitatively or quantitatively extreme. Arthur's expectation of the incubation time for spoiled fish can range from being identical to the actual events, or wildly incorrect. Our investigation of temporal congruity in Chapter 5 simply examined ratings by participants who expected either a short delay or a long delay and were presented with either a short or a long delay. However, this problem intersects an interesting line research in the covariation assessment literature that examines the influence of outliers on judgements of covariation (e.g., Jennings, Amabile, & Ross, 1982a). How are participants' judgements affected by events that are either *extremely* congruent or incongruent with their expectations? The problem of temporal congruity lends itself well to investigating this question.

6.0.4 Shortcomings and Future Directions

The seven experiments presented in Chapters 2 and 3 tested whether the fact that participants were "predicting" or "diagnosing" influenced the relative assessment of the events. While the manipulations presented in Chapter 3 demonstrated that we could "push around" their sensitivity to the direction and nature of the causal relationship, the effects were not as large as I would have liked. In each of these experiments, participants seemed to regard one effect in light of another to some extent resulting in a considerable cue-interaction effect in the 2E-1C scenario. While the integrative test question had a moderate influence on reducing this cue-interaction effect, the results did not reach the level of the predictions made by causal-model theory. It seems plausible that using a different type of cover story for the 2E-1C scenario would result in data that are more consistent with the predictions of causal-model theory. As indicated in Chapter 3, Waldmann and Holyoak (1997) insisted on the necessity of making the causal relationship in the instructions and materials unmistakable. They argued that if there is any ambiguity in participants' interpretation of the event relations, then one cannot accurately measure the influence of their causal interpretation on learning. In Experiment 1 of Chapter 3 we responded to this criticism by continually reminding participants about the nature and direction of the causal relationship after every 8 trials and immediately before making a causal judgement. The manipulation, however, did not influence participants' sensitivity to the causal relationship resulting in a substantial cue-interaction effect in both the 2C-1E and 2E-1C causal scenarios. One could argue, however, that the reversibility of chemical-bacteria cover story may still be influencing their causal assessments. Therefore, one may wish to investigate the same experimental design, but use materials for the 2E-1C scenario that cannot possibly be interpreted any other way than two effects resulting from a common cause. For example, two symptoms (headache and stomach pains) may result from eating a particular food. It is unlikely that participants would regard the headache and stomach pains as *causing* the consumption of the food. Though I would speculate that if enough trials were presented, then the causal description of the events may eventually lose their significance resulting in associative-type behaviour (cue-interaction between effects).

The integrative test question described in Chapter 3 was designed so participants would *use*

the causal model for some particular purpose and regard all of the causal information presented in the entire experiment. While this manipulation was moderately successful in influencing participants' sensitivity to the causal description of the events, future work in this area might consider other manipulations that influence the *use* of the events. For example, participants in the 2C-1E and 2E-1C conditions are presented with a series of four cue combinations: AB, A∼B, ∼AB, and ∼A∼B. One could tell the participants that scientists have since discovered a fifth bacterial strain in the human digestive system. They are interested in the extent to which Chemical Alpha and Chemical Beta influence this new bacterial strain (the extent to which this new strain influences the production of Chemical Alpha and Chemical Beta). However, unlike the previous strains, experimental trials on this new bacterial strain are extremely expensive. If you were given $100 to spend on the four trial types (AB, A∼B, ∼AB, and ∼A∼B), how much of the $100 would you spend on each in order to best assess the influence of each chemical on the bacteria (the influence of the bacteria on each chemical). If participants are behaving in accord with causal-model theory, then they should allocate more funds to the A∼B and ∼AB trial types because they should conditionalise, or hold one cause constant while evaluating the other. Participants in the 2E-1C condition, on the other hand, should allocate funds equally to the four types (or at least allocate less funds to the A∼B and ∼AB trial types), because they should regard the influence of the cause on each effect independently. Other types of manipulations aimed at participants use of the causal models would likely result in data that are more consistent with causal-model theory.

The work presented in this dissertation has reached an intersection where basic associative processes meet more cognitive phenomena. Each chapter has described circumstances in which expectations, or high-level cognitive phenomena influence basic formations of association. These investigations coincide nicely with the neighbouring field of covariation assessment.

While continuing to investigate many of the topics presented in this dissertation, I hope to examine the role of expectations and the treatment of outliers in the area of covariation estimation. As indicated in the previous section, by extending the dichotomous variables that I used in the experiments presented here, to form continuous variables, I gain the possibility of investigating extreme observations. How is our ability to profit from data that we receive affected by expectations about what that relationship should be? For example, does a causal model reduce our sensitivity to data that contradicts that expectation? There are two underlying issues in examining people's attitudes toward data analysis and the treatment of outliers depending on whether they are congruent or incongruent with their expectations. They can be ignored and regarded as a different process to be rejected — noise to be seen through. This would likely be the case if one's expectations were incongruent with the events presented and, in particular, with the *extreme* events presented. Alternatively, one can focus on the outliers and regard them as entirely representative of the underlying process. Outliers are defined in terms of surprise, attention, memory, and theoretical expectations. It would be interesting to see whether one's interpretation of outliers varies according to one's expectations.

6.1 Conclusion

The thirteen experiments in this dissertation were designed to assess the circumstances under which knowledge of causal direction and temporal delay influence judgements of causality. In Chapters 2 and 3 the extent that participants causal assessments were influenced by knowledge of causal asymmetry depended on how and when the judgements were obtained. Participants were more sensitive to the causal structure of the events in their ratings than their trial-by-trial prediction responses, on earlier rather than later trials, and when asked to provide an integrative causal rating. Emphasising the direction and nature of the causal relationship, and the wording of the test question had no influence on participants' sensitivity to causal asymmetry. The results from the experiments in Chapters 2 and 3 are best explained by the joint contribution of high and low-level processing whereby causal judgements are influenced both by causal knowledge and by the associative nature of the events. The results from Chapter 4 suggest that participants' ratings track the conditional rather than uncondtional contingencies as predicted by the conditional ΔP account as well as the Rescorla-Wagner model at asymptote. The three related experiments in the postscript of Chapter 4 suggest that participants tend to rate the influence of each cause conditional on the absence of the other cause. Finally, Chapter 5 suggests that knowledge of temporal delay modulates causal judgements.

References

Allan, L. G. (1980). A note on measurement of contingency between two binary variables in judgment tasks. *Bulletin of the Psychonomic Society, 15*, 147–149.

Allan, L. G. (1993). Human contingency judgments: Rule based or associative? *Psychological Bulletin, 114*, 435–448.

Allan, L. G. (2003). Assessing power PC. *Learning and Behavior, 31*, 192-204.

Allan, L. G., Balsam, P., Church, R. M., & Terrace, H. (2002). John gibbon: Obituary. *American Psychologist, 57*, 436–437.

Allan, L. G., & Church, R. M. (2002). Introduction. *Learning and Motivation: Special Issue to honor the work of John Gibbon, 33*, 311–334.

Allan, L. G., & Jenkins, H. M. (1983). The effect of representations of binary variables on judgments of influence. *Learning and Motivation, 14*, 381–405.

Allan, L. G., Tangen, J. M., Wood, R., & Shah, T. (in press). Temporal contiguity and contingency judgments: A Pavlovian analogue. *Integrative Physiological and Behavior Science.*

Baker, A. G., Mercier, P., Vallée-Tourangeau, F., Frank, R., & Pan, M. (1993). Selective association and causality judgments: Presence of a strong causal factor may reduce judgments of a weaker one. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*, 414–432.

Baker, A. G., Murphy, R. A., & Vallée-Tourangeau, F. (1996). Associative and normative models of cause. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The psychology of learning and motivation* (pp. 1–45). San Diego: Academic Press.

Buehner, M. J., & May, J. (2002). Knowledge mediates the timeframe of covariation assessment in human causal induction. *Thinking & Reasoning, 8*, 269–295.

Buehner, M. J., & May, J. (in press). Rethinking temporal contiguity and the judgment of causality: Effects of prior knowledge, experience, and reinforcement procedure. *Quarterly Journal of Experimental Psychology. A, Human Experimental Psychology.*

Chapman, G. B., & Robbins, S. J. (1990). Cue interaction in human contingency judgment. *Memory & Cognition, 18*, 537–545.

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review, 104*, 367–405.

Cheng, P. W., & Holyoak, K. J. (1995). Comparative approaches to cognitive science. In H. L. Roitblat & J. A. Meyer (Eds.), (pp. 271–302). Cambridge, MA: MIT Press.

Cheng, P. W., & Novick, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology, 58*, 545–567.

Cheng, P. W., & Novick, L. R. (1992). Covariation in natural causal induction. *Psychological Review, 99*, 365–382.

Cobos, P. L., López, F. J., Caño, A., Almaraz, J., & Shanks, D. R. (2002). Mechanisms of predictive and diagnostic causal induction. *Journal of Experimental Psychology: Animal Behavior Processes, 28*, 331-346.

Collins, D. J., & Shanks, D. R. (2002). Momentary and integrative response strategies in causal judgment. *Memory & Cognition, 30*, 1138–1147.

Collins, D. J., & Shanks, D. R. (in press). Momentary and integrative response strategies in causal judgment. *Memory & Cognition.*

Danks, D. (in press). Equilibria of the Rescorla-Wagner model. *Journal of Mathematical Psychology.*

Dickinson, A. (2001). Causal learning: An associative analysis. *Quarterly Journal of Experimental Psychology, 54B*, 3-25.

Dickinson, A., & Burke, J. (1996). Within-compound associations mediate the retrospective revaluation of causality judgements. *Quarterly Journal of Experimental Psychology, 49B*, 60-80.

Dickinson, A., & Shanks, D. R. (1985). Animal learning and human causality judgment. In L. G. Nilson & T. Archer (Eds.), *Perspectives on learning and memory* (pp. 167–191). Hillsdale: Lawrence Erlbaum Associates.

Dickinson, A., Shanks, D. R., & Evendon, J. (1984). Judgment of act-outcome contingency: The role of selective attribution. *Quarterly Journal of Experimental Psychology, 36A*, 29–50.

Doolittle, M. H. (1888). Association ratios. *Bull. Philos. Soc. Washington, 10*, 83–96.

Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart.* New York: Oxford University Press.

Hagmayer, Y., & Waldmann, M. R. (2000). Simulating causal models: The way to structural sensitivity. In L. R. Gleitman & A. K. Joshi (Eds.), *Proceedings of the twenty-second annual conference of the cognitive science society* (Vol. 82, pp. 214–219). Mahwah, NJ: Erlbaum.

Hagmayer, Y., & Waldmann, M. R. (2003). How temporal assumptions influence causal judgments. *Memory & Cognition, 30*, 1128–1137.

Hall, G., & Pearce, J. M. (1979). Latent inhibition of a CS during CS-US pairings. *Journal of Experimental Psychology: Animal Behavior Processes, 5*, 31–42.

Hausman, D. M. (1993). Why don't effects explain their causes? *Synthese*, 227-244.

Howell, D. C. (1997). *Statistical methods for psychology* (4th ed.). Boston: Duxbury Press.

Jennings, D. L., Amabile, T., & Ross, L. (1982a). Informal covariation assessment: Data-based versus theory-based judgments. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 211–230). Cambridge: Cambridge University Press.

Jennings, D. L., Amabile, T., & Ross, L. (1982b). Judgment under uncertainty: Heuristics and biases. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 3–20). Cambridge: Cambridge University Press.

Kamin, L. J. (1968). "Attention-like" processes in classical conditioning. In M. R. Jones (Ed.), *Miami symposium on the production of behavior: Aversive stimulation* (pp. 9–33). Miami, Florida: University of Miami Press.

Kamin, L. J. (1969a). Predictability, surprise, attention, and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior* (pp. 279–296). New York: Appleton-Century-Crofts.

Kamin, L. J. (1969b). Selective association and conditioning. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental issues in associative learning* (pp. 42–64). Halifax, Nova Scotia, Canada: Dalhousie University Press.

Lober, K., & Shanks, D. R. (2000). Is causal induction based on causal power? Critique of Cheng (1997). *Psychological Review, 107*, 195-212.

López, F. J., Shanks, D. R., Almaraz, J., & Fernandez, P. (1998). Effects of trial order on contingency judgments: A comparison of associative and probabilistic contrast accounts. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24*, 672–694.

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcment. *Psychological Review, 82*, 276–298.

Matute, H., Arcediano, F., & Miller, R. R. (1996). Test question modulates cue competition between causes and between effects. *Journal of Experimental Psychology: Learning Memory, and Cognition, 22*, 182–196.

Matute, H., Vegas, S., & De Marez, P. (2002). Flexible use of recent information in causal and predictive judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*, 714-725.

Mehta, R. R. (2000). *Contrasting associative and statistical theories of contingency judgments.* Unpublished doctoral dissertation, McGill University.

Miller, R. R., & Barnet, R. C. (1993). The role of time in elementary associations. *Current Directions in Psychological Science, 2*, 106–111.

Nisbett, R., & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment.* New Jersey: Prentice-Hall.

Pearce, J. M. (1987). A model of stimulus generalization for Pavlovian conditioning. *Psychological Review, 94*, 61–73.

Popper, K. (1972). *Conjectures and refutations.* London: Routledge and Kegan Paul.

Price, P. C., & Yates, F. (1993). Judgmental overshadowing: Further evidence of cue interaction in contingency judgment. *Memory & Cognition, 21*, 561–572.

Price, P. C., & Yates, F. (1995). Associative and rule-based accounts of cue interaction in contingency judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 1639–1655.

Reed, P. (1992). Effect of a signalled delay between an action and outcome on human judgment of causality. *Quarterly Journal of Experimental Psychology, 44B*, 81–100.

Reed, P. (1996). No evidence for blocking in human judgments of causality by stimuli presented during an outcome delay. *Learning and Motivation, 27*, 317–333.

Reichenbach, H. (1956). *The direction of time.* Berkeley and Los Angeles: University of California Press.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokosy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.

Sadeghi, H. (2003). *Cue interaction and judgements of causality.* Unpublished honours thesis, McMaster University.

Savastano, H. I., & Miller, R. R. (1998). Time as content in Pavlovian conditioning. *Behavioural Processes, 44*, 147–162.

Savastano, H. I., Yin, H., Barnet, R. C., & Miller, R. R. (1998). Temporal coding in Pavlovian conditioning: Hall-Pearce negative transfer. *Quarterly Journal of Experimental Psychology, 51B*, 139–153.

Shanks, D. R. (1985). Forward and backward blocking in human contingency judgement. *Quarterly Journal of Experimental Psychology. B, Comparative and Physiological Psychology, 37B*, 1–21.

Shanks, D. R. (1989). Selectional processes in causality judgment. *Memory & Cognition, 17*, 27–34.

Shanks, D. R. (1993). Human instrumental learning: A critical review of data and theory. *British Journal of Psychology, 84*, 319-354.

Shanks, D. R., & Dickinson, A. (1987). Associative accounts of causality judgment. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 21, pp. 229–261). San Diego: Academic Press.

Shanks, D. R., & Dickinson, A. (1991). Instrumental judgment and performance under variations in action-outcome contingency and contiguity. *Memory & Cognition, 19*, 353–360.

Shanks, D. R., Holyoak, K. J., & Medin, D. L. (1996). *The psychology of learning and motivation: Causal learning.* San Diego, CA: Academic Press.

Shanks, D. R., & López, F. J. (1996). Causal order does not affect cue selection in human associative learning. *Memory & Cognition, 24*, 511–522.

Shanks, D. R., Pearson, S. M., & Dickinson, A. (1989). Temporal contiguity and the judgement of causality by human subjects. *Quarterly Journal of Experimental Psychology, 41B*, 139–159.

Siegel, S., & Allan, L. G. (1996). The widespread influence of the Rescorla-Wagner model. *Psychological Bulletin and Review, 3*, 314–321.

Spellman, B. A. (1996a). Acting as intuitive scientists: Contingency judgments are made while controlling for alternative potential causes. *Psychological Science, 7*, 337–342.

Spellman, B. A. (1996b). Conditionalizing causality. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The psychology of learning and motivation: Vol. 34. Causal learning* (pp. 167–206). San Diego, CA: Academic Press.

Spellman, B. A., Price, C. M., & Logan, J. M. (2001). How two causes are different from one: The use of (un)conditional information in Simpson's paradox. *Memory & Cognition, 29*, 193-208.

Tangen, J. M., & Allan, L. G. (2003). The relative effect of cue interaction. *Quarterly Journal of Experimental Psychology. B, Comparative and Physiological Psychology, 56*, 279-300.

Tangen, J. M., & Allan, L. G. (submitted). Cue-interaction and judgments of causality. *Memory & Cognition.*

Tangen, J. M., Allan, L. G., & Sadeghi, H. (invited). Assessing (in)sensitivity to causal asymmetry: A matter of degree. In A. J. Wills (Ed.), *New directions in human associative learning.* Hillsdale, NJ: Erlbaum.

Vallée-Tourangeau, F., Baker, A. G., & Mercier, P. (1994). Discounting in causality and covariation judgements. *Quarterly Journal of Experimental Psychology. B, Comparative and Physiological Psychology, 47B*, 151–171.

Vallée-Tourangeau, F., Murphy, R. A., & Baker, A. G. (1998). Casual induction in the presence of a perfect negative cue: Contrasting predictions from associative and statistical models. *Quarterly Journal of Experimental Psychology, 51B*, 173-191.

Vallée-Tourangeau, F., Murphy, R. A., & Drew, S. (1997). Causal judgments that violate the predictions of the power PC theory of causal induction. In M. G. Shafto & P. Langley (Eds.), *Proceedings of the nineteenth annual conference of the Cognitive Science Society* (p. 775-780). Hillsdale NJ: Erlbaum.

Vallée-Tourangeau, F., Murphy, R. A., Drew, S., & Baker, A. (1998). Judging the importance of constant and variable candidate causes: A test of the power PC theory. *Quarterly Journal of Experimental Psychology, 51A*, 65-84.

Van Hamme, L. J., Kao, S. F., & Wasserman, E. A. (1993). Judging interevent relations: From cause to effect and from effect to cause. *Memory & Cognition, 21*, 802–808.

Van Hamme, L. J., & Wasserman, E. A. (1993). Cue competition in causality judgements: The role of manner of information presentation. *Bulletin of the Psychonomic Society, 31*, 457-460.

Van Hamme, L. J., & Wasserman, E. A. (1994). Cue competition in causality judgments: The role of nonpresentation of compound stimulus elements. *Learning and Motivation, 25*, 127-151.

Wagner, A. R. (1981). SOP: A model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5–47). Hillsdale, NJ: Erlbaum.

Wagner, A. R., Logan, F. A., Haberlandt, K., & Price, T. (1968). Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology, 76*, 171–180.

Waldmann, M. R. (2000). Competition among causes but not effects in predictive and diagnostic learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*, 53–76.

Waldmann, M. R. (2001). Predictive versus diagnostic causal learning: Evidence from an overshadowing paradigm. *Psychonomic Bulletin and Review, 8*, 600–608.

Waldmann, M. R., & Hagmayer, Y. (1999). How categories shape causality. In M. Hahn & S. C. Stoness (Eds.), *Proceedings of the twenty-first annual conference of the cognitive science society* (Vol. 82, pp. 761–766). Mahwah, NJ: Erlbaum.

Waldmann, M. R., & Hagmayer, Y. (2001). Estimating causal strength: The role of structural knowledge and processing effort. *Cognition, 82*, 27–58.

Waldmann, M. R., & Holyoak, K. J. (1990). Can causal induction be reduced to associative learning? In *Proceedings of the twelfth annual conference of the cognitive science society* (pp. 190–197). Hillsdale, NJ: Erlbaum.

Waldmann, M. R., & Holyoak, K. J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General, 121*, 222–236.

Waldmann, M. R., & Holyoak, K. J. (1997). Determining whether causal order affects cue selection in human contingency learning: Comments on Shanks and Lopez (1996). *Memory & Cognition, 25*, 125–134.

Waldmann, M. R., Holyoak, K. J., & Fratianne, A. (1995). Causal models and the acquisition of category structure. *Journal of Experimental Psychology: General, 124*, 181–206.

Wasserman, E. A. (1990). Attribution of causality to common and distinctive elements of compound stimuli. *Psychological Science, 1*, 298–302.

Wasserman, E. A., & Neunaber, D. J. (1986). College students responding to and rating of contingency relations: The role of temporal contiguity. *Journal of the Experimental Analysis of Behavior, 46*, 15–35.

Williams, D. A. (1996). A comparative analysis of negative contingency learning in humans and nonhumans. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The psychology of learning and motivation* (pp. 89–125). San Diego: Academic Press.

Appendix A

Causal Model Prompts

A.0.1 2C-1E

1. You are assessing how well each chemical causes the strain of bacteria to survive.
2. The two chemicals are causes and the bacteria is the effect.
3. Each of the two chemicals may have a positive, negative, or no influence on the survival of the bacteria.
4. Your job is to evaluate the extent to which each chemical aids in or interferes with the survival of the bacterial strain.
5. Consider the causal influence of each chemical on the survival of the bacteria.
6. The goal is to predict whether the bacterial strain will survive given the presence or absence of the chemicals.
7. You are trying to determine the influence of each chemical on the survival of the bacteria.
8. Think about the affect each chemical has on the bacterial strain.
9. Consider the causal strength of each chemical.
10. Try to judge the effectiveness of each chemical by its influence on the bacteria.
11. Keep in mind that each chemical affects the bacterial strain.
12. Decide the degree to which the bacterial strain depends on the addition of each chemical.
13. Try to keep track of what happens to the bacterial strain when one, both, or neither chemical was added.
14. Each chemical may cause the bacteria to survive more often, less often, or may have no influence on the strain's survival.

15. Evaluate the extent to which each chemical caused the strain of bacteria to survive.
16. Try to ascertain the causal influence of each chemical on the bacteria.

A.0.2 2E-1C

1. You are assessing how well the bacterial strain causes each of the two chemicals to be produced.
2. The bacteria is the cause and the two chemicals are the effects.
3. The bacteria may have a positive, negative, or no influence on the production of each chemical.
4. Your job is to evaluate the extent to which the bacterial strain aids in or interferes with the production of each chemical.
5. Consider the causal influence of the bacteria on the production of each chemical.
6. The goal is to diagnose whether the bacteria were added to the digestive system given the presence or absence of the chemicals.
7. You are trying to determine the influence of of the bacteria on the production of the two chemicals.
8. Think about the affect the bacterial strain has on each chemical.
9. Consider the causal strength of the bacterial strain.
10. Try to judge the effectiveness of the bacteria by its influence on each chemical.
11. Keep in mind that the bacterial strain affects each chemical.
12. Decide the degree to which each chemical depends on the addition of the bacterial strain.
13. Try to keep track of whether one, both, or neither chemical was produced when the bacterial strain was added.
14. The bacteria may cause the production of each chemical more often, less often, or may have no influence on whether the chemical is produced.
15. Evaluate the extent to which the strain of bacteria caused each chemical to be produced.
16. Try to ascertain the causal influence of the bacteria on each chemical.

Appendix B

Instructions for Chapter 4

Imagine that scientists have recently discovered three strains of bacteria that exist in the mammalian digestive system. The scientists are studying whether certain pairs of chemicals affect the bacteria's survival. For each strain, the scientists are testing whether the chemicals aid in, interfere with, or have no effect on the strain's survival.

To do this, a strain of bacteria was first placed in culture (petri dishes). After that,

1. one chemical (e.g., chemical A)
2. the other chemical (e.g., chemical B)
3. both chemicals (e.g., chemicals A and B); or
4. neither chemical

was added to the bacterial culture. A few hours later, the scientists verified whether or not the bacterial sample survived.

At the time of these experiments, it was not known what effect each chemical might have had on the bacteria. On one hand, a chemical might make the bacteria *more* likely to survive, therefore the sample would be less likely to survive without the chemical. This would be an example of a chemical having a *positive* effect on the bacteria's survival.

Alternatively, a chemical might make it *less* likely that that bacteria will survive, therefore the sample would be more likely to survive without the chemical. This would be an example of a chemical having a *negative* effect on the bacteria's survival.

Finally, a chemical might have no systematic effect on the bacteria's survival. That is, it could be that the chemical neither aids in, nor interferes with, the bacteria's survival.

To assess these possibilities, the scientists investigated what happened when one or two chemicals were added to the bacteria sample. They also tested what happened on control trials in which no chemicals were added. A comparison of what happened on these trials allowed the scientists to assess whether a chemical had a *positive* effect, a *negative* effect or no effect on the bacteria's survival.

To ensure that the results were reliable, the following measures were taken: - The scientists verified that the chemicals used did not interact. That is, when mixed, they neither

neutralized each other nor formed a more potent compound.
- Similar concentrations of chemicals were used in each experiment.
- The age and concentration of the bacteria were similar in all conditions.
- The optimal conditions for the bacteria's survival were first established. That is, each strain's optimal temperature, pH, lighting, and nutrients were verified prior to the beginning of the experiments and these conditions were consistently used.
- The experiments were conducted under sterile conditions. That is, the cultures were first checked to ensure that they were not contaminated. As well, the scientists ensured that the samples were not exposed to contaminants in the air.
- The scientists verified that the test used to classify the bacteria as surviving or not was reliable. In previous studies, it was found to yield results equally reliable to those obtained from counting the bacteria before and after adding chemicals whose actions were known to bacterial samples.

You will be presented with the results from this study. On each trial, you will be told whether one chemical, the other chemical, both chemicals, or no chemicals were added to the bacterial sample. You will then decide whether you think the bacterial sample will survive. Use the computer mouse to click the appropriate button. Click "Survived" if you think the bacteria survived, or "Did Not Survive" if you think the bacteria did not survive. After clicking the button, you will be told whether your guess is correct or incorrect. If a chemical was not added to the bacterial sample, then a picture of the chemical will appear in gray. Similarly, if the bacterial strain did not survive, then a picture of the bacteria will appear in gray.

When doing the task, try to keep track of what happened when one or both chemicals were added, as well as what happened when neither chemical was added. However, do not write down this information.

You will be presented with the results from three experiments. At the end of each experiment, you will be asked to rate how strongly each chemical affects the bacteria's survival. You will rate the strength on a scale ranging from -100 to 100. Remember that a *positive* number means that you think that the chemical has a positive effect on the bacteria's survival. That is, the bacteria are more likely to survive if the chemical is added than if it is not added. A *negative* number means that you think that the chemical has a negative effect on the bacteria's survival. That is, the bacteria are less likely to survive if the chemical is added than if it is not added. And zero means that the chemical does not systematically aid in nor interfere with the bacteria's survival.

The number you enter indicates how strongly positive or negative you think is the chemical's effect on the bacteria. 100 means that the chemical has a very strong positive effect, while a rating such as 50 means that the chemical has a moderately positive effect on the bacteria's survival. Similarly, -100 means that the chemical has a strong negative effect on the bacteria's survival, while a rating such as -25 means that the chemical has a weak negative effect on the bacteria's survival.

Appendix C

Rescorla-Wagner = Conditional ΔP

In the one-phase blocking task there are eight types of trials corresponding to the eight cells of the 4×2 matrix. On each trial there is an equation for each presented cue. The equations are shown below for each cell of the matrix. In the matrix, $\lambda = 1$ on outcome present trials and $\lambda = 0$ on outcome absent trials. As in Figure 4.2, the frequencies of the cue-outcome combinations are represented by the letters a, b, ...

	O	~O
ABX	**a** $\Delta V_X = \alpha_X\beta[1 - (V_X + V_A + V_B)]$ $\Delta V_A = \alpha_A\beta[1 - (V_X + V_A + V_B)]$ $\Delta V_B = \alpha_B\beta[1 - (V_X + V_A + V_B)]$	**b** $\Delta V_X = \alpha_X\beta[0 - (V_X + V_A + V_B)]$ $\Delta V_A = \alpha_A\beta[0 - (V_X + V_A + V_B)]$ $\Delta V_B = \alpha_B\beta[0 - (V_X + V_A + V_B)]$
AX	**c** $\Delta V_X = \alpha_X\beta[1 - (V_X + V_A)]$ $\Delta V_A = \alpha_A\beta[1 - (V_X + V_A)]$	**d** $\Delta V_X = \alpha_X\beta[0 - (V_X + V_A)]$ $\Delta V_A = \alpha_A\beta[0 - (V_X + V_A)]$
BX	**e** $\Delta V_X = \alpha_X\beta[1 - (V_X + V_B)]$ $\Delta V_B = \alpha_B\beta[1 - (V_X + V_B)]$	**f** $\Delta V_X = \alpha_X\beta[0 - (V_X + V_B)]$ $\Delta V_B = \alpha_B\beta[0 - (V_X + V_B)]$
X	**g** $\Delta V_X = \alpha_X\beta[1 - (V_X)]$	**h** $\Delta V_X = \alpha_X\beta[0 - (V_X)]$

$$\text{mean}\,\Delta V_A = \frac{a\alpha_A\beta[1-(V_X+V_A+V_B)] + b\alpha_A\beta[0-(V_X+V_A+V_B)] + c\alpha_A\beta[1-(V_X+V_A)] + d\alpha_A\beta[0-(V_X+V_A)]}{a+b+c+d} \tag{C.1}$$

$$\text{mean}\,\Delta V_B = \frac{a\alpha_B\beta[1-(V_X+V_A+V_B)] + b\alpha_B\beta[0-(V_X+V_A+V_B)] + e\alpha_B\beta[1-(V_X+V_B)] + f\alpha_B\beta[0-(V_X+V_B)]}{a+b+e+f} \tag{C.2}$$

$$\text{mean}\,\Delta V_X = \frac{a\alpha_X\beta[1-(V_X+V_A+V_B)] + b\alpha_X\beta[0-(V_X+V_A+V_B)] + c\alpha_X\beta[1-(V_X+V_A)] + d\alpha_X\beta[0-(V_X+V_A)]}{a+b+c+d+e+f+g+h} + \frac{e\alpha_X\beta[1-(V_X+V_A+V_B)] + b\alpha_X\beta[0-(V_X+V_B)] + g\alpha_X\beta[1-(V_X)] + h\alpha_X\beta[0-(V_X)]}{a+b+c+d+e+f+g+h} \tag{C.3}$$

Equations (C.1)-(C.3) simplify to

$$\text{mean}\,\Delta V_A = a\alpha_A\beta\frac{a+c-V_A(a+b+c+d)-V_B(a+b)-V_X(a+b+c+d)}{a+b+c+d} \tag{C.4}$$

$$\text{mean}\,\Delta V_B = a\alpha_B\beta\frac{a+e-V_A(a+b)-V_B(a+b+e+f)-V_X(a+b+e+f)}{a+b+e+f} \tag{C.5}$$

$$\text{mean}\,\Delta V_X = a\alpha_X\beta\frac{a+c+e+g-V_A(a+b+c+d)-V_B(a+b+e+f)-V_X(a+b+c+d+e+f+g+h)}{a+b+c+d+e+f+g+h} \tag{C.6}$$

At asymptote $\text{mean}\,\Delta V_A = 0$, $\text{mean}\,\Delta V_B = 0$, $\text{mean}\,\Delta V_X = 0$, and Equations (C.4)-(C.6) simplify to

$$a+c-V_A(a+b+c+d)-V_B(a+b)-V_X(a+b+c+d)=0 \tag{C.7}$$

$$a+e-V_A(a+b)-V_B(a+b+e+f)-V_X(a+b+e+f)=0 \tag{C.8}$$

$$a+c+e+g-V_A(a+b+c+d)-V_B(a+b+e+f)-V_X(a+b+c+d+e+f+g+h)=0 \tag{C.9}$$

For Property 1,

$$b = g \tag{C.10a}$$

$$d = e \tag{C.10b}$$

$$f = c \tag{C.10c}$$

$$h = a \tag{C.10d}$$

For Property 4,

$$(a + b) = (c + d) = (e + f) = (g + h) \tag{C.11}$$

Substituting Equations (C.10a), (C.10b), and (C.11) in Equation (C.9),

$$2V_X = 1 - V_A - V_B \tag{C.12}$$

Substituting Equation (C.11) in Equation (C.7),

$$a + c - (c + d)(2V_A + V_B + 2V_X) = 0 \tag{C.13}$$

Substituting Equation (C.12) in Equation (C.13),

$$V_A = \frac{a + c}{a + b} - 1$$

$$V_A = \frac{a + c - c - d}{a + b} = \frac{a - d}{a + b} = \Delta P_A \tag{C.14}$$

Substituting Equation (C.11) in Equation (C.8),

$$a + d - (c + d)(V_A + 2V_B + 2V_X) = 0 \tag{C.15}$$

Substituting Equation (C.12) in Equation (C.15),

$$V_B = \frac{a + d}{a + b} - 1$$

$$V_B = \frac{a + d - c - d}{a + b} = \frac{a - c}{a + b} = \Delta P_B \tag{C.16}$$

Appendix D

Instructions for Chapter 5

D.1 Instructions for Experiment 1

Imagine that scientists have recently discovered four strains of bacteria that exist in the mammalian digestive system. The scientists are studying whether certain chemicals affect the bacteria's survival. For each strain, the scientists are testing whether the chemicals aid in, interfere with, or have no effect on the strain's survival.

To do this, a strain of bacteria was first placed in culture (petri dishes). After that, a chemical was either added or not added to the bacterial culture. A few hours later, the scientists verified whether or not the bacterial sample survived.

At the time of these experiments, it was not known what effect each chemical might have had on the bacteria. On one hand, a chemical might make the bacteria MORE likely to survive, therefore the sample would be less likely to survive without the chemical. This would be an example of a chemical having a POSITIVE effect on the bacteria's survival.

Alternatively, a chemical might make it LESS likely that that bacteria will survive, therefore the sample would be more likely to survive without the chemical. This would be an example of a chemical having a NEGATIVE effect on the bacteria's survival.

Finally, a chemical might have NO systematic effect on the bacteria's survival. That is, it could be that the chemical neither aids in, nor interferes with, the bacteria's survival.

To assess these possibilities, the scientists investigated what happened when a chemical added to the bacteria sample. They also tested what happened on control trials in which no chemical was added. A comparison of what happened on these trials allowed the scientists to assess whether a chemical had a POSITIVE effect, a NEGATIVE effect or NO effect on the bacteria's survival.

To ensure that the results were reliable, the following measures were taken: - Similar concentrations of chemicals were used in each experiment.
- The age and concentration of the bacteria were similar in all conditions.
- The optimal conditions for the bacteria's survival were first established. That is, each strain's optimal temperature, pH, lighting, and nutrients were verified prior to the beginning of the experiments and these conditions were consistently used.

- The experiments were conducted under sterile conditions. That is, the cultures were first checked to ensure that they were not contaminated. As well, the scientists ensured that the samples were not exposed to contaminants in the air.
- The scientists verified that the test used to classify the bacteria as surviving or not was reliable. In previous studies, it was found to yield results equally reliable to those obtained from counting the bacteria before and after adding chemicals whose actions were known to bacterial samples.

You will be presented with the results from this study. On each trial, you will be told whether the chemical was added to the bacterial sample. You will then be told whether the bacteria survived or did not survive. If the chemical was not added to the bacterial sample, then a picture of the chemical will appear in gray. Similarly, if the bacterial strain did not survive, then a picture of the bacteria will appear in gray.

When doing the task, try to keep track of what happened when the chemical was added, as well as what happened when the chemical was not added. However, do not write down this information.

You will be presented with the results from four experiments. At the end of each experiment, you will be asked to rate how strongly the chemical affects the bacteria's survival. You will rate the strength on a scale ranging from -100 to +100. Remember that a POSITIVE number means that you think that the chemical has a postive effect on the bacteria's survival. That is, the bacteria are more likely to survive if the chemical is added than if it is not added. A NEGATIVE number means that you think that the chemical has a negative effect on the bacteria's survival. That is, the bacteria are less likely to survive if the chemical is added than if it is not added. And zero means that the chemical does not systematically aid in nor interfere with the bacteria's survival.

The number you enter indicates how strongly postive or negative you think is the chemical's effect on the bacteria. +100 means that the chemical has a very strong positive effect, while a rating such as +50 means that the chemical has a moderately positive effect on the bacteria's survival. Similarly, -100 means that the chemical has a strong negative effect on the bacteria's survival, while a rating such as -25 means that the chemical has a weak negative effect on the bacteria's survival.

D.2 Instructions for Experiment 2

D.2.1 Immediate Instructions - Remote-Control Detonator

The army is testing a number of new remote control detonators. Imagine that you are a military officer at an army training site, and that you are observing the results of a series of tests concerned with the effectiveness of these remote control detonators. Your task is to decide whether triggering the FIRE button on the remote control detonator is effective in setting off the explosion of a mine in the training range. The detonator, when fired, emits a radio signal, so that clicking of the FIRE button should produce an *immediate* explosion.

The detonators are still in the experimental phase. You will see that triggering the FIRE

button does not always result in an explosion of the mine. You will also see that some of the mines explode spontaneously even when the FIRE button is not clicked.

The test of a particular detonator consists on a number of trials. On each trial, you will see whether the FIRE button was activated or not, and then whether the mine exploded or not. Remember, that because of the radio signal, the mine should explode *immediately* after the clicking of the FIRE button.

At the end of a series of trials you will be asked to rate the relationship between clicking the FIRE button and the explosion of a mine. You will rate the strength of the relationship on a scale ranging from 0 to 100. A rating of zero means that clicking the FIRE button had no effect on causing the explosion; that is, the explosion was just as likely to be spontaneous as caused by the FIRE button. A rating of 100 means that the FIRE button was a perfect cause of the explosion; that is, the explosion was never spontaneous and was always caused by the FIRE button. Between the extremes, your rating reflects the increase in explosions of the mine caused by the clicking of the FIRE button.

During the course of the experiment, you will evaluate four different remote detonators. Some of these detonators might be more effective than others.

D.2.2 Delay Instructions - Grenade Launcher

The army is testing a number of new grenade launchers. Imagine that you are a military officer at an army training site, and that you are observing the results of a series of tests concerned with the effectiveness of these grenade launchers. Your task is to decide whether triggering the FIRE button on the launcher is effective in setting off the explosion of a mine in the training range. The grenade launcher, when fired, sends shells into the training range. The shells have to travel from the launcher to the mine site, so there should be a few seconds *delay* between the clicking of the FIRE button and the mine explosion.

The launchers are still in the experimental phase. You will see that triggering the FIRE button does not always result in an explosion of the mine. You will also see that some of the mines explode spontaneously even when the FIRE button is not clicked.

The test of a particular launcher consists on a number of trials. On each trial, you will see whether the FIRE button was activated or not, and then whether the mine exploded or not. Remember, that because of the distance of the launcher from the mine site, there should be a *delay* between clicking the FIRE button and the mine explosion.

At the end of a series of trials you will be asked to rate the relationship between clicking the FIRE button and the explosion of a mine. You will rate the strength of the relationship on a scale ranging from 0 to 100. A rating of zero means that clicking the FIRE button had no effect on causing the explosion; that is, the explosion was just as likely to be spontaneous as caused by the FIRE button. A rating of 100 means that the FIRE button was a perfect cause of the explosion; that is, the explosion was never spontaneous and was always caused by the FIRE button. Between the extremes, your rating reflects the increase in explosions of the mine caused by the clicking of the FIRE button.

During the course of the experiment, you will evaluate four different grenade launchers. Some of these launchers might be more effective than others.